THE LONDON SCHOOL OF ECONOMICS AND POLITICAL SCIENCE

ESSAYS ON FINANCIAL MARKETS AND BUSINESS CYCLES

HANS FABIAN WINKLER

A thesis submitted to the
Department of Economics of the London School of Economics
for the degree of Doctor of Philosophy

May 2015

Hans Fabian Winkler: *Essays on Financial Markets and Business Cycles.* A thesis submitted for the degree of Doctor of Philosophy in Economics at the London School of Economics and Political Science, © May 2015.

DECLARATION

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 36,987 words.

STATEMENT OF CONJOINT WORK

I confirm that Chapter 3 was jointly co-authored with Dr. Stéphane Moyen and Dr. Nikolai Stähler at the Deutsche Bundesbank. The opinions expressed in the chapter do not necessarily reflect the views of the Deutsche Bundesbank, the Eurosystem or its staff. I wrote and analysed the simplified model, implemented the solution of the Ramsey problems, and wrote the chapter.

ABSTRACT

This thesis contains three essays on the linkages between financial markets and business cycles.

The first chapter introduces a method to embed learning about asset prices (relying on past observation to predict future prices) into business cycle models in a way that retains a maximum of rationality and parsimony. This method is applied to a real business cycle model and a search model of unemployment. In the RBC model, learning about stock prices leads to counterfactual correlations between consumption, employment and investment. By contrast, the search model augmented by learning can generate realistic business cycle fluctuations. The volatility of unemployment in the data can be replicated without the need to rely on a high degree of wage rigidity.

The second chapter examines the implications of a learning-based asset pricing theory for a model of firm financial frictions. Learning greatly improves asset price properties such as return volatility and predictability. In combination with financial frictions, a powerful feedback loop emerges between beliefs, stock prices and real activity, leading to substantial amplification of shocks. The model-implied subjective expectations are found to be consistent with patterns of forecast error predictability in survey data. A reaction of monetary policy to asset prices stabilises expectations and substantially improves welfare, which is not the case under rational expectations.

The third chapter is concerned with the inefficiencies caused by incomplete national and international financial markets. Specifically, it examines the optimal design of an unemployment insurance scheme that operates across multiple countries in the presence of such inefficiencies. Using a two-country business cycle model with labour market search frictions, it is found that a supranational unemployment insurance scheme can be used to achieve transfers across countries without changing unemployment levels; and that the optimal unemployment insurance policy prescribes a countercyclical replacement rate due in the presence of cross-country transfers.

ACKNOWLEDGEMENTS

During my years as a PhD student, I have received support and advice from many people.

First and foremost, I thank my advisor Wouter den Haan for the many hours he has spent discussing, encouraging and criticising my work. He always found time to answer questions and read through drafts, paying attention to the big picture and the technical details. I learned as much from his knowledge of economics as I did from his work attitude and intellectual integrity.

I am also indebted to Albert Marcet who took the time to give advice on my thesis in its final year and allowed me to spend several weeks at IAE-CSIC in Barcelona. The discussions with him have greatly helped me sharpen my ideas.

Stéphane Moyen and Nikolai Stähler have been great coauthors on Chapter 3 of this thesis. I have learned much from them about the interplay between economic research and policy.

This thesis has also benefited greatly from many comments and suggestions. In particular, I would like to thank Johannes Boehm, Huaizhi Chen, Elena Gerko, Keyu Jin, Max Klimm, Pascal Michaillat, Rachel Ngai, Markus Riegler, Kevin Sheedy, and Shengxing Zhang for fruitful discussions.

No matter how much I thank my parents, it will be little compared to the love and support I received from them. They always encouraged me in my plans even in difficult times.

Finally, I thank my fiancée Shing Lei Koh for sharing all better and worse moments of life with me. Through her love, I am able to see further and clearer.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

1 THE EFFECT OF LEARNING ABOUT FIRM VALUE IN RBC AND LABOUR SEARCH MODELS

1.1 INTRODUCTION

Asset prices are hard to predict, but they are not just random walks. At least since Fama and French (1988) wrote about the negative correlation of stock prices with future returns, we know that asset prices exhibit some degree of mean reversion. But we know much less about why this is so. Proponents of the efficient market hypothesis tend to say that high asset prices are a manifestation of low required returns (e.g. Cochrane, 2011). When investors only require low returns to hold assets (for example because of an increased desire to save or increased risk appetite), asset prices rise today and expectedly fall back in the future. However, an alternative view holds that high asset prices are instead a manifestation of high expected returns (e.g. Adam et al., 2013). When investors are optimistic about the future, they buy assets at high prices in the expectation of even further price appreciation. Mean reversion in prices is then unexpected by investors. Survey data on actual investor expectations seem to favour this alternative view (Greenwood and Shleifer, 2014). However, it is not without trouble, because it goes against the rational expectations hypothesis: investors must repeatedly and systematically ignore available evidence on return predictability, instead forming some sort of extrapolative expectations or learning behaviour.

In this chapter, I take the extrapolative expectations approach to asset pricing at face value. Suppose that agents in the economy really do form expectations about asset prices through a learning mechanism. How does this affect our understanding of business cycles?

Most of the models used to understand business cycles have rational expectations as a fundamental building block. They are also not very good at asset pricing since the marginal investor has time-separable preferences with low coefficients of relative risk aversion. By adding asset price learning, could we improve the asset pricing properties of those models and at the same time learn something new about real fluctuations?

I analyse two of the most commonly used business cycle models: a RBC model and a labour search and matching model. For the RBC model, I find that learning does improve some asset price properties, but also leads to counterfactual negative comovement between consumption, output and stock prices. The reason for this is that stock prices have no allocative role in the RBC model other than determining the tradeoff between consumption and investment. A stock price boom will lead to a corresponding boom in consumption at the expense of investment and also employment through a wealth effect. This comovement problem is in fact well known from the literature on news shocks (Jaimovich and Rebelo, 2009).

In the search and matching model of unemployment, the comovement problem can be overcome. Here, learning about firm value leads to additional amplification of productivity shocks to all main macroeconomic aggregates. The business cycle properties are as good as under rational expectations, while the asset price properties improve. Moreover, the model with learning substantially reduces the need to rely on wage rigidity to match the observed volatility of unemployment fluctuations. It has been argued previously that the standard search model can only produce realistic unemployment volatility when firm profits are very volatile (Shimer, 2005; Hagedorn and Manovskii, 2008). This chapter shows that the problem is rather the low volatility of firm *value*.

Abandoning rational expectations in a business cycle model is not without problems. Rational expectations have the advantage that the modeller does not have to worry about how to specify an expectations formation process. Without them, many degrees of freedom need to be filled, and the existing attempts in the learning literature to do so (e.g. Milani, 2011) often become intransparent. To address this problem, I build on the concept of internal rationality developed by Adam and Marcet (2011). An internally rational equilibrium is one in which agents do make choices by maximising an intertemporal objective function with coherent and time-consistent beliefs about the variables affecting their decisions, but their beliefs do not necessarily have to coincide with the actual probability distributions of those variables. This concept still leaves many degrees of freedom: In a business cycle model, for example, households could still entertain arbitrary beliefs about future wages and firm profits, and firms could entertain arbitrary beliefs about required returns (discount factors). I therefore develop a refinement of their equilibrium concept, which I call "*conditonally model-consistent expectations*" (CMCE).

With CMCE, agents can entertain arbitrary beliefs about one relative price in the economy (which will be the price of stocks in this chapter), but their beliefs about any other variable must be consistent with optimal choices of other agents and market clearing in all but two markets. Spelling out a belief system for stock prices and then imposing conditionally model-consistent expectations is all that is needed to obtain a unique dynamic equilibrium for the models studied in this chapter. What's more, when agents with CMCE evaluate their forecast errors, they find that their forecasting rules cannot be improved upon as long as they hold on to the original subjective belief about stock prices. In this sense, this equilibrium refinement represents a minimal departure from rational expectations.

The remainder of this chapter is structured as follows. Section 1.2 reviews the concept of internal rationality and introduces the refinement of conditionally model-consistent expectations. Section 1.3 illustrates this concept by applying it to a simple endowment economy with stock price learning in two different ways. Only one of them is suitable for a solution via perturbation methods, the other has to be solved with global methods because of the non-linearities involved. Section 1.4 applies the analysis to the RBC model, while Section 1.5 deals with the search and matching model. Section 1.6 concludes.

1.2 CONDITIONALLY MODEL-CONSISTENT EXPECTATIONS

The purpose of this section is to define the equilibrium concept used in this chapter. I build on the concept of "internal rationality" introduced by Adam and Marcet (2011). An internally rational equilibrium relaxes rational expectations and allows for the construction of an equilibrium when agents have arbitrary beliefs about variables that they take as external and given in their decision problem. While intuitive, internal rationality leads to many degrees of freedom when specifying beliefs. I develop a particular specification of expectations that allows for arbitrary beliefs in one, and only one, relative price that agents can learn about, and otherwise stay as close as possible to rational expectations. I call this "conditionally model-consistent expectations". It will be shown that rational expectations are nested in this concept as a special case, and that agents endowed with conditionally model-consistent expectations make the best possible forecasts for all other model variables, conditional on the relative price they are learning about.

1.2.1 *General formulation*

I consider a macroeconomic model of fairly general form. Time is discrete at periods $t = 0, 1, 2, \ldots$. The economy consists of a number of agents, indexed by $i \in I$. Agents can trade various commodities in competitive spot markets, such as consumption goods, labour, financial assets etc. These goods are indexed by $j \in J$. The economy is stochastic through a number of exogenous variables, collected in the joint stochastic process $\mathbf{u} = (u_t)_{t=0}^{\infty}$, defined over a probability space $(\Omega_u, \sigma_u, \mathcal{P}_u)$. In particular, the probability measure $\mathcal{P}_u$ describes the actual and objective distribution of shocks.

We are interested in determining the endogenous model variables. These fall into two categories. The first category are choices made by agents as a result of some optimisation problem. Let the choices made by agent i be $\mathbf{y}_i = (y_{it})_{t=0}^{\infty}$. The second category are market-clearing prices. Each good $j \in J$ trades at the price p_{jt} in period t. All endogenous model variables are collected in the stochastic process $\mathbf{y} = (y_t)_{t=0}^{\infty} = \left((y_{it})_{i \in I}, (p_{jt})_{j \in J}\right)_{t=0}^{\infty}$.

Choices made by agents are the outcome of optimisation problems. I will describe them in a fairly abstract form. Each agent $i \in I$ observes a set of model variables relevant to his decision problem. These "inputs" into the decision problem are denoted by $\mathbf{x}_i = (x_{it})_{t=0}^{\infty}$. The agent chooses $\mathbf{y}_i^*$ by solving an optimisation problem of the following general form:

$$\mathbf{y}_i^* \in \arg\max_{\mathbf{y}_i \in G_i} \int F_i(\mathbf{y}_i, \mathbf{x}_i) \mathcal{P}_i(d\mathbf{x}_i)$$

Here, the functional F_i is a path-wise objective function. For a household, this might be the present discounted sum of period utility evaluated in expectation at time zero. The probability measure $\mathcal{P}_i$ is the *subjective belief* of agent i about the stochastic distribution of the decision inputs $\mathbf{x}_i$.[1] Importantly, this belief $\mathcal{P}_i$ can (for now) be arbitrary and is not restricted to satisfy rational expectations. The set G_i describes the admissible solutions. This set can incorporate all kinds of constraints on the optimisation problem, such as budget constraints or boundary conditions. Also, G_i needs to restrict the solutions to be adapted with respect to $\mathbf{x}_i$. This simply means that in any period t, the solution y_{it} can only depend on the information contained in $(x_{i0}, \ldots, x_{it})$ but not on future realisations $x_{it+1}, x_{it+2}, \ldots$. The solution to this problem is really a mapping from inputs to choices, and I will therefore denote it by a function $\mathbf{x}_i \mapsto \mathbf{y}_i^*(\mathbf{x}_i)$.

As for the determination of prices, I require that equilibrium prices clear markets. In the most general form, market clearing of commodity $j \in J$ in every period will require that $H_j(\mathbf{y}) = 0$ for some function H_j.

We now have the notation in place to define internal rationality, rational expectations, and conditionally model-consistent expectations.

Agents are internally rational, in the words of Adam and Marcet (2011), when they "make fully optimal decisions given a well-defined system of subjective probability beliefs about payoff relevant variables that are beyond their control or 'external', including prices". These beliefs, however, need not coincide with the actual probability distribution of external variables in equilibrium. Using the notation above, this can be formalised as follows.

Definition 1.1. A collection of beliefs $(\mathcal{P}_i)_{i \in I}$ and a mapping $g : \mathbf{u} \mapsto \mathbf{y}$ is called an *internally rational equilibrium* if

1. $g(\mathbf{u}) = (y_t)_{t=0}^{\infty}$ is an adapted process with respect to $\mathbf{u}$, so that y_t depends only on $(u_0, \ldots, u_t)$;

1 For example, if $x_{it} \in \mathbb{R}$ then $\mathbf{x}_i \in \mathbb{R}^{\mathbb{N}}$, and $\mathcal{P}_i$ is a probability measure defined on $\mathbb{R}^{\mathbb{N}}$ with its corresponding Borel sigma-algebra.

2. for every agent $i \in I$, there exists a solution $\mathbf{x}_i \mapsto \mathbf{y}_i^* (\mathbf{x}_i)$ to the agent's optimisation problem with belief $\mathcal{P}_i$ such that $g_{\mathbf{y}_i} (\mathbf{u}) = \mathbf{y}_i^* (g_{\mathbf{x}_i} (\mathbf{u}))$ $\mathcal{P}_u$-almost surely;

3. the allocations defined by g clear all markets: For all $j \in J$: $H_j (g (\mathbf{u})) = 0$ $\mathcal{P}_u$-almost surely.

Condition 1 states that equilibrium outcomes at time t can only depend on information available in the economy up to time t. Condition 2 requires that the equilibrium outcomes are consistent with each agent's optimisation problem, solved under the subjective belief $\mathcal{P}_i$. Condition 3 imposes market clearing.

This sounds similar to the familiar and standard definition of a rational expectations equilibrium, but there is an important difference: The subjective beliefs $\mathcal{P}_i$ and the actual distribution of equilibrium outcomes (defined by $\mathcal{P}_u$ and g) need not coincide. A rational expectations equilibrium is a special case of an internally rational equilibrium.

Definition 1.2. A *rational expectations equilibrium* is an internally rational equilibrium with the additional property that, for every agent i, beliefs about the distribution of $\mathbf{x}_i$ under $\mathcal{P}_i$ coincide with equilibrium outcomes under $\mathcal{P}_u$: $\mathbf{x}_i \mid \mathcal{P}_i \stackrel{d}{=} g_{\mathbf{x}_i} (\mathbf{u}) \mid \mathcal{P}_u$.

Rational expectations (RE) have the advantage that they leave no degrees of freedom for specifying agents' beliefs about their environment. They are the solution to an elegant fixed-point problem: Beliefs have to coincide with the equilibrium distribution of model outcomes, and model outcomes have to coincide with optimal choices given beliefs. But even when the rational expectations equilibrium of a model is unique, the number of internally rational equilibria is usually very large, as there are no constraints on the subjective beliefs $\mathcal{P}_i$. In Adam et al. (2014), there is only one agent (the representative investor) and only one external decision input (the stock price), so spelling out a belief system is still manageable. But particularly for larger business cycle models, filling the degrees of freedom is a difficult task. To allow for *some* departure from rational expectations while keeping the degrees of freedom as small as possible, I introduce conditionally model-consistent expectations (CMCE).

Under RE, beliefs $\mathcal{P}_i$ are pinned down by the actual distribution of exogenous variables $\mathcal{P}_u$ and the equilibrium mapping $g : \mathbf{u} \mapsto (\mathbf{y}, \mathbf{u})$ to all model variables. With just internal rationality, the beliefs $\mathcal{P}_i$ can in principle be completely arbitrary. Under CMCE, we start by picking one competitively traded good $j_0 \in J$, and we can specify an arbitrary belief for its relative price. Formally, define a distribution $\tilde{\mathcal{P}}$ over the joint process $(\mathbf{u}, \mathbf{p}_{j_0})$ of exogenous variables and the relative price of the j_0-good as the basis of the belief system. Conditionally model-consistent expectations then pin down the expectations of all other model variables and a mapping $h : \left(\mathbf{u}, \mathbf{p}_{j_0}\right) \mapsto \mathbf{y}$. This mapping has to satisfy a number of criteria. I will first spell out the formal definition.

Definition 1.3. An internally rational equilibrium $\left(g, (\mathcal{P}_i)_{i \in I}\right)$ has *conditionally model-consistent expectations with respect to* $\mathbf{p}_{j_0}$ if there exists a probability measure $\tilde{\mathcal{P}}$ over $\left(\mathbf{u}, \mathbf{p}_{j_0}\right)$ and a mapping $h : \left(\mathbf{u}, \mathbf{p}_{j_0}\right) \mapsto \mathbf{y}$ such that:

1. for every agent $i \in I$, the subjective belief is defined by $\tilde{\mathcal{P}}$ and h: $\mathbf{x}_i \mid \mathcal{P}_i \sim h_{\mathbf{x}_i}\left(\mathbf{u}, \mathbf{p}_{j_0}\right) \mid \mathcal{P}$;
2. $h\left(\mathbf{u}, \mathbf{p}_{j_0}\right)$ is an adapted process with respect to $\left(\mathbf{u}, \mathbf{p}_{j_0}\right)$ under $\tilde{\mathcal{P}}$, so that y_t depends only on $\left(u_0, p_{j_0 0}, \ldots, u_t, p_{j_0 t}\right)$;
3. for every agent $i \in I$, there exists a solution $\mathbf{x}_i \mapsto \mathbf{y}_i^*\left(\mathbf{x}_i\right)$ to his optimisation problem with belief $\mathcal{P}_i$ such that: $\mathbf{y}_i^*\left(h_{\mathbf{x}_i}\left(\mathbf{u}, \mathbf{p}_{j_0}\right)\right) = h_{\mathbf{y}_i}\left(\mathbf{u}, \mathbf{p}_{j_0}\right)$ $\tilde{\mathcal{P}}$-almost surely;
4. the allocations defined by h clear all markets *except* the market for j_0 and one other market: there exists a $j_1 \in J$ such that $H_j\left(h\left(\mathbf{u}, \mathbf{p}_{j_0}\right)\right) = 0$ $\tilde{\mathcal{P}}$-almost surely for all $j \in J \setminus \{j_0, j_1\}$;
5. beliefs about exogenous variables are correct: $\mathbf{u} \mid \tilde{\mathcal{P}} \overset{d}{=} \mathbf{u} \mid \mathcal{P}_u$;
6. subjective beliefs are consistent with equilibrium outcomes on the equilibrium path: $h\left(\mathbf{u}, g_{\mathbf{p}_{j_0}}\left(\mathbf{u}\right)\right) = g\left(\mathbf{u}\right)$ $\mathcal{P}_u$-almost surely.

Condition 1 states that the subjective beliefs $\mathcal{P}_i$ entering the decision problems of the agents are defined by some joint belief about exogenous variables $\mathbf{u}$ and the price $\mathbf{p}_{j_0}$ and a deterministic mapping from those to all other model variables. Condition 2 states that agents believe that the state of the economy at time t can only depend on the realisation of exogenous variables and the stock price up to time t. Condition 3 states that beliefs are consistent with the choices of agents. In practice, this means that whenever the choices of one agent enter the decision problem of another agent, the beliefs of the other agent about those choices have to be consistent with what the first agent would choose, for any realisation of fundamentals $\mathbf{u}$ and prices $\mathbf{p}_{j_0}$. Condition 4 says that beliefs have to be consistent with market clearing. In practice, this means that whenever a competitive market price enters the decision problem of an agent, the beliefs of that agent need to be market-clearing prices. This has to hold for any market $j \in J$ with the exception of j_0 and one other market j_1.[2] Condition

2 Why not impose that only the market for j_0 does not clear? Since the mapping h is consistent with optimal choices, Walras law applies to it and the clearing of all markets except one implies clearing of all markets already. But a mapping h that is consistent with total market clearing and optimal choices under beliefs induced by h is necessarily a rational expectations mapping (see the previous definition). Therefore, to allow from departures from rational expectations, two markets have to be in disequilibrium in the subjective belief system.

5 imposes that agents know the true distribution of the exogenous variables $\mathbf{u}$. Finally, Condition 6 states that agents' beliefs (defined by the function h) coincide with model outcomes (defined by the function g) when stock prices are in equilibrium.

Intuitively, CMCE impose a similar fixed-point logic onto beliefs as RE, but exclude the relative price $\mathbf{p}_{j_0}$ from it. Choices have to be optimal given beliefs, and beliefs have to be consistent with choices and clearing of all markets *excluding* the market for j_0 and one more market. Removing this market clearing condition leaves one degree of freedom which can be filled by a subjective belief for $\mathbf{p}_{j_0}$. This belief and the other equilibrium conditions pin down the subjective policy function h. The actual model outcomes g are then found by imposing the remaining market clearing condition, which imposes an equilibrium price in good j_0.

This definition retains many elements of a rational expectations equilibrium. RE are also model-consistent expectations, and as the name suggests, they are a special case of CMCE.

Proposition 1.4. *Any rational expectations equilibrium $\left(g, (\mathcal{P}_i)_{i\in I}\right)$ has conditionally model-consistent expectations with respect to $\mathbf{p}_{j_0}$ for any $j_0 \in J$.*

Proof. Define the subjective probability measure by $\tilde{\mathcal{P}}\left(\left(\mathbf{u}, \mathbf{p}_{j_0}\right) \in \mathcal{A}\right) = \mathcal{P}_u\left(\varphi_{j_0}\left(g\left(\mathbf{u}\right)\right) \in \mathcal{A}\right)$. Define the mapping h by $h\left(\mathbf{u}, \mathbf{p}_{j_0}\right) = g\left(\mathbf{u}\right)$. By the properties of a rational expectations equilibrium, we have $\mathbf{x}_i \mid \mathcal{P}_i \stackrel{d}{=} g_{\mathbf{x}_i}\left(\mathbf{u}\right) \mid \mathcal{P}_u$. From the construction of $\tilde{\mathcal{P}}$ and h it follows that $\mathbf{x}_i \mid \mathcal{P}_i \sim h_{\mathbf{x}_i}\left(\mathbf{u}, \mathbf{p}_{j_0}\right) \mid \tilde{\mathcal{P}}$. This establishes the first property in the definition of CMCE. The second, third and fourth property follow from the definition of internal rationality and the fact that $h\left(\mathbf{u}, \mathbf{p}_{j_0}\right) \mid \tilde{\mathcal{P}} \sim g\left(\mathbf{u}\right) \mid \mathcal{P}_u$. The fifth property is established by the fact that $\tilde{\mathcal{P}}\left(\mathbf{u} \in \mathcal{B}\right) = \tilde{\mathcal{P}}\left(\left(\mathbf{u}, \mathbf{p}_{j_0}\right) \in \mathcal{B} \times \Omega_{\mathbf{p}_{j_0}}\right) = \mathcal{P}_u\left(\varphi_{j_0}\left(g\left(\mathbf{u}\right)\right) \in \mathcal{B} \times \Omega_{\mathbf{p}_{j_0}}\right) = \mathcal{P}_u\left(\mathbf{u} \in \mathcal{B}\right)$ (where $\Omega_{\mathbf{p}_{j_0}}$ is the range of $\mathbf{p}_{j_0}$). Finally the sixth property, $h\left(\mathbf{u}, g_{\mathbf{p}_{j_0}}\left(\mathbf{u}\right)\right) = g\left(\mathbf{u}\right)$, follows directly from the construction of h. □

Another way to characterise conditionally model-consistent expectations is through forecast errors. Would an agent with such expectations have reason to revise his predictions after observing the realised model outcomes? The answer is that, conditional on observations of the prices $\mathbf{p}_{j_0}$, he would not, as his forecasts are perfectly consistent with his observations.

Proposition 1.5. *For an internally rational equilibrium that has conditionally model-consistent expectations with respect to $\mathbf{p}_{j_0}$, let $\left(p_{j_0 t}\right)_{t=0}^{\infty} = g_{\mathbf{p}_{j_0}}\left(\mathbf{u}\right)$ and $\left(x_{it}\right)_{t=0}^{\infty} = g_{\mathbf{x}_i}\left(\mathbf{u}\right)$ be equilibrium model outcomes. Then for all agents $i \in I$ and $t, s \geq 0$, it holds that*

$$\forall i \in I\, \forall t, s \geq 0: \mathbb{E}_{t-s}^{\mathcal{P}_i}\left[x_{it} \mid u_{t-s}, p_{j_0 t-s}, \ldots, u_t, p_{j_0 t}\right] - x_{it} = 0\, \mathcal{P}_u\text{-a.s.}$$

Proof. Using the defining properties, we obtain:

$$\begin{aligned}
& \mathbb{E}_{t-s}^{\mathcal{P}_i}\left[x_{it} \mid u_{t-s}, p_{j_o t-s}, \ldots, u_t, p_{j_o t}\right] \\
= \; & \mathbb{E}_{t-s}^{\mathcal{P}}\left[h_{x_{it}}\left(\mathbf{u}, \mathbf{p}_{j0}\right) \mid u_{t-s}, p_{j_o t-s}, \ldots, u_t, p_{j_o t}\right] \\
= \; & \mathbb{E}^{\mathcal{P}}\left[h_{x_{it}}\left(\mathbf{u}, \mathbf{p}_{j0}\right) \mid u_0, p_{j_o}, \ldots, u_t, p_{j_o t}\right] \\
= \; & h_{x_{it}}\left(\mathbf{u}, \mathbf{p}_{j0}\right) \\
= \; & h_{x_{it}}\left(\mathbf{u}, g_{\mathbf{p}_{j0}}\left(\mathbf{u}\right)\right) \\
= \; & g_{x_{it}}\left(\mathbf{u}\right) \\
= \; & x_{it}.
\end{aligned}$$

□

The above proposition effectively says that if the agent maintains his belief $\tilde{\mathcal{P}}$ about the exogenous variables $\mathbf{u}$ and prices $\mathbf{p}_{j0}$, his forecasts cannot be further improved. They are the best possible model forecasts given $\tilde{\mathcal{P}}$, since they produce zero forecast errors when conditioned on $\left(\mathbf{u}, \mathbf{p}_{j0}\right)$. This is a strong property of conditionally model-consistent expectations.

1.2.2 *Recursive formulation*

In practice, macroeconomic models often take a recursive form, which greatly facilitates their solution. It is therefore instructive to rephrase the notion of CMCE in a recursive form as well.

A macroeconomic model can often be brought into the following recursive form. The exogenous shocks are a white noise process $(u_t)_{t=0}^{\infty}$ where $u_t \in \mathbb{R}^p$ are independent and identically distributed, with mean zero, variance $\sigma^2\Sigma_u$ and distribution F. The endogenous model variables form a stochastic process $(y_t)_{t=0}^{\infty}$ where $y_t \in \mathbb{R}^n$. A solution of the model takes the form of a policy function:

$$y_t = g\left(y_{t-1}, u_t, \sigma\right)$$

Under rational expectations, this policy function is the solution to a system of n recursive equilibrium conditions as follows:

$$\begin{aligned}
0 &= \mathbb{E}_t^{\mathcal{P}_u}\left[f\left(y_{t+1}, y_t, y_{t-1}, u_t\right)\right] \\
&= \int_{\mathbb{R}^p} f\left(g\left(g\left(y, u, \sigma\right), u', \sigma\right), g\left(y, u, \sigma\right), y, u\right) F\left(du'\right)
\end{aligned}$$

where $f : \mathbb{R}^{n\times n\times n\times p} \to \mathbb{R}^n$. These n conditions contain the optimality conditions of all agents as well as the market-clearing conditions.[3] Note that the optimality conditions for any one agent in this formulation are evaluated using the true distribution of shocks and all other equilibrium conditions: That is the essence of rational expectations.

3 When there are J competitive markets, the equation system will include $J - 1$ market clearing conditions, as one is redundant by Walras law.

This formulation can be extended to solve for internally rational equilibria with conditionally model-consistent expectations. Before solving the *actual* law of motion of the economy, one has to solve for its *perceived* law of motion under subjective beliefs. Subjective beliefs are determined by a belief $\mathcal{P}$ about the evolution of the price p_{j_0t}, and a mapping h from exogenous shocks u_t and the price p_{j_0t}, where h is consistent with optimal choice and market clearing in all but two markets.

We can freely specify a belief about p_{j_0t} by a system of equations of the form:

$$0 = \psi\left(y_t, y_{t-1}, u_t, z_t\right)$$

where $\psi : \mathbb{R}^{n\times n\times p\times 1} \to \mathbb{R}$ and $(z_t)_{t=0}^{\infty}$ is an additional white noise process. The $z_t \in \mathbb{R}$ are independent and identically distributed, with mean zero, variance $\sigma^2\Sigma_z$ and distribution function G. In the applications of this chapter, they will correspond to agents' subjective forecast errors on the price p_{j_0t}. The goal is to get a recursive representation of the perceived law of motion under CMCE of the form:

$$y_t = h\left(y_{t-1}, u_t, z_t, \sigma\right)$$

Under conditionally model-consistent expectations, the mapping h must be consistent with all the equilibrium conditions in f except for the clearing condition of market j_0. Collect these $n-1$ equilibrium conditions in the function $f_{-j_0} : \mathbb{R}^{n\times n\times n\times p} \to \mathbb{R}^{n-1}$ and the j_0-market clearing condition in the function $\phi : \mathbb{R}^{n\times n\times n\times p} \to \mathbb{R}$. A recursive system of equations describing the subjective belief is given by:

$$\begin{aligned}
0 &= \begin{pmatrix} \psi\left(y_t, y_{t-1}, u_t, z_t\right) \\ \mathbb{E}_t^{\mathcal{P}}\left[f_{-j_0}\left(y_{t+1}, y_t, y_{t-1}, u_t\right)\right] \end{pmatrix} \\
&= \int_{\mathbb{R}^p}\int_{\mathbb{R}} \begin{pmatrix} \psi\left(h\left(y,u,z,\sigma\right), y, u, z\right) \\ f_{-j_0}\begin{pmatrix} h\left(h\left(y,u,z,\sigma\right), u', z', \sigma\right), \\ h\left(y,u,z,\sigma\right), y, u \end{pmatrix} \end{pmatrix} G\left(dz'\right) F\left(du'\right)
\end{aligned}$$

This is a system of n equations which can be solved for the subjective policy function h.

In a second step, one has to find the *actual* law of motion g of the economy. This is done by imposing clearing of market j_0. Market clearing will determine an equilibrium path for the price p_{j_0t} or, alternatively, for the subjective forecast errors z_t. It is convenient to solve for a policy function that determines z_t as follows:

$$z_t = r\left(y_{t-1}, u_t, \sigma\right)$$

This can be solved using the remaining market clearing condition ϕ:

$$\begin{aligned}
0 &= \mathbb{E}_t^{\mathcal{P}}\phi\left(y_{t+1}, y_t, y_{t-1}, u_t\right) \\
&= \int_{\mathbb{R}^p}\int_{\mathbb{R}} \phi\begin{pmatrix} h\left(h\left(y,u,r\left(y,u,\sigma\right),\sigma\right), u', z', \sigma\right), \\ h\left(y,u,r\left(y,u,\sigma\right),\sigma\right), y, u \end{pmatrix} G\left(dz'\right) F\left(du'\right)
\end{aligned}$$

The actual law of motion can now be written as:

$$y_t = g\left(y_{t-1}, u_t, \sigma\right) = h\left(y_{t-1}, u_t, r\left(y_{t-1}, u_t, \sigma\right), \sigma\right)$$

The difference between what agents think will happen and what actually happens lies in the forecast error z_t of the price $p_{j_0 t}$. Agents think that z_t is a random variable with a distribution, but actually it is a deterministic function of the state variables and the other exogenous shocks. Agents never understand that z_t does not follow the distribution they believe in. But they understand perfectly well how the economy behaves for a given value of z_t.

Standard numerical methods can be applied to solve for an internally rational equilibrium with CMCE in this form. In Appendix A.1, I describe how standard perturbation methods can be used to compute second order approximations to the policy functions in practice.

1.3 ENDOWMENT ECONOMY

This section applies the concept of conditionally model-consistent expectations (CMCE) to a simple endowment economy without production and with learning about stock prices as in Adam et al. (2014). This serves primarily to make the concept more familiar to the reader. I offer two ways to introduce CMCE. The first way leads to learning dynamics identical to those in Adam et al.. However, it is unsuitable for low-order approximation by perturbation methods, which makes it problematic to embed this approach into larger business cycle models. The second way is more suitable for perturbation methods and will be used for the business cycle models in this chapter.

1.3.1 *Model setup*

Time is discrete at periods $t = 0, 1, 2, \ldots$. There are two consumption goods. The first consumption good is a "regular" consumption good. Flows of this good come from two assets. The first asset delivers a constant endowment $W_t = \bar{W}$ which is paid every period. The second is a Lucas tree that delivers a stochastic endowment A_t, where

$$\log A_t = \rho \log A_{t-1} + \varepsilon_t \tag{1.1}$$

and $\varepsilon_t \sim \mathcal{N}\left(0, \sigma_\varepsilon^2\right)$ is a white noise process. Both assets are in unitary supply. There is no production, so in equilibrium total consumption has to equal $A_t + \bar{W}$.

I call the second consumption good the "zero-good". This good is in zero net supply, hence the choice of name. Regular consumption goods and zero-goods cannot be transformed into each other. The presence of this good will only matter in the absence of rational expectations, as will become clear shortly.

The economy is populated by a representative household, an investment agency and a firm. The household receives the constant endowment $\bar{W}$ and also receives profits π_t from the investment agency. The household only values regular consumption goods. It maximises utility as follows:

$$\max_{(C_t)_{t=0}^{\infty}} \mathbb{E}^{\mathcal{P}} \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\gamma}}{1-\gamma}$$

$$\text{s.t. } C_t \leq W_t + \pi_t.$$

The investment agency acts on behalf of the household and can purchase shares in the firm at price P_t, each share entitling it to a dividend payment D_t. It solves the following problem:

$$\max_{\left(S_t,\pi_t,C_t^0\right)_{t=0}^{\infty}} \mathbb{E}^{\mathcal{P}} \sum_{t=0}^{\infty} Q_{0,t} \left(\pi_t + C_t^0\right)$$

$$\begin{aligned} \text{s.t. } \pi_t + C_t^0 + P_t S_t &\leq (P_t + D_t)\, S_{t-1} \\ S_t &\in [0, \bar{S}] \end{aligned}$$

where the household's discount factor is $Q_{t,s} = \beta^{s-t} \left(C_s/C_t\right)^{-\gamma}$. Stock holdings of the agency are bounded from below by zero (no short-selling) and from above by a threshold $\bar{S} > 1$.[4] The agency's objective includes the present discounted value of profits paid to the household, discounted at the household rate, and also consumption goods C_t^0 of the zero-good purchased in the market and transferred to the household.

This objective function is not consistent with the household's utility function, as the zero good C_t^0 provides no value to the household. Note though that in any equilibrium, it has to be the case that $C_t^0 = 0$ since the zero good is in zero net supply. The zero-good will only play a role if the investment agency expects non-zero future values of C_t^0 in the future in the absence of rational expectations. In any event, the first order condition of the investment agency is given by:

$$S_t = 0 \quad \text{if} \quad P_t > \mathbb{E}_t^{\mathcal{P}} \left[Q_{t,t+1} \left(P_{t+1} + D_{t+1}\right)\right], \tag{1.2}$$

$$S_t \in [0, \bar{S}] \quad \text{if} \quad P_t = \mathbb{E}_t^{\mathcal{P}} \left[Q_{t,t+1} \left(P_{t+1} + D_{t+1}\right)\right], \tag{1.3}$$

$$S_t = \bar{S} \quad \text{if} \quad P_t < \mathbb{E}_t^{\mathcal{P}} \left[Q_{t,t+1} \left(P_{t+1} + D_{t+1}\right)\right]. \tag{1.4}$$

Finally, the firm is simply administering the Lucas tree. It can not issue new shares or buy back old ones, and the supply of shares is fixed at $S_t = 1$. The firm maximises dividend payments to the household, using the latter's discount factor:

$$\max_{(D_t)_{t=0}^{\infty}} \mathbb{E}^{\mathcal{P}} \sum_{t=0}^{\infty} Q_{0,t} D_t$$

4 This guarantees the existence of a solution with finite stock holdings under arbitrary beliefs. In equilibrium these constraints are never binding.

$$\text{s.t. } D_t \leq A_t.$$

Clearly, the firm will choose $D_t = A_t$.

The purpose of this model is to discuss the properties of the asset price P_t under different informational assumptions. The model setup might at first seem peculiar. Why do we need to separate out a firm and an investment agency in an endowment economy? Why do we need to introduce a "zero-good"? Indeed, under rational expectations, all these modelling choices are superfluous. But the distinction will serve to clarify the concept of conditionally model-consistent expectations.

1.3.2 *Rational expectations equilibrium*

Any equilibrium with market clearing must satisfy $S_t = 1$ and $C_t^0 = 0$. Since both firm and household want to exhaust their budget constraint, it also has to be the case that $D_t = A_t$ and $C_t = \bar{W} + A_t$. This completely determines allocations. The only interesting question is which price process P_t prevails in equilibrium. Under rational expectations, the answer is easy. Using the equilibrium conditions, the price P_t solves

$$\begin{aligned} P_t &= \mathbb{E}_t \left[Q_{t,t+1} \left(P_{t+1} + D_{t+1} \right) \right] \\ &= \mathbb{E}_t \left[\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \left(P_{t+1} + D_{t+1} \right) \right] \\ &= \mathbb{E}_t \left[\beta \left(\frac{W_{t+1} + D_{t+1}}{W_t + D_t} \right)^{-\gamma} \left(P_{t+1} + D_{t+1} \right) \right] \\ &= \mathbb{E}_t \left[\beta \left(\frac{\bar{W} + A_{t+1}}{\bar{W} + A_t} \right)^{-\gamma} \left(P_{t+1} + A_{t+1} \right) \right] \\ &= \mathbb{E} \left[\sum_{s=1}^{\infty} \beta^s \left(\frac{\bar{W} + A_{t+s}}{\bar{W} + A_t} \right)^{-\gamma} A_{t+s} \right]. \end{aligned} \tag{1.5}$$

The first equality must always hold as long as the investment agency takes decisions optimally and the stock market is in equilibrium. But all the remaining equalities make use of rational expectations. The second equality holds because the investment agency knows the equilibrium distribution of the discount factor. The third equality holds because the household expects $S_t = 1$ to be his optimal choice at all times (since prices must be such that $S_t = 1$ and the household's expectation of prices is consistent with this). The fourth equality holds because the household knows perfectly well that the endowment asset delivers a constant flow of consumption goods $W_t = \bar{W}$, and he also knows that the firm will optimally choose $D_t = A_t$. The fifth equality holds by forward iteration of the optimality condition and repeated application of the previous equalities.

1.3.3 *Learning equilibrium*

An internally rational equilibrium needs to specify beliefs about variables external to every agent's decision problem, respectively. There are three agents, the household, the investment agency and the firm. The inputs $\mathbf{x}_i$ into the household's decision problems are the evolution of the endowment W_t and the profits π_t received by the investment agency. The household's choices $\mathbf{y}_i$ consist in consumption C_t. The inputs $\mathbf{x}_i$ into the investment agency's decision problem are the dividends D_t received from the firm; the stock price P_t; and the household discount factor $Q_{t,t+1}$. The agency's choices $\mathbf{y}_i$ are stockholdings S_t, profits π_t and transfers C_t^0 of zero-goods. The inputs $\mathbf{x}_i$ into the firm's problem are the stochastic endowment A_t, and the decision $\mathbf{y}_i$ is the dividend payment D_t.

I could specify an internally rational equilibrium for a variety of beliefs about all these decision inputs, so the number of degrees of freedom is rather large even in this simple example. The concept of conditionally model-consistent expectations (CMCE) greatly reduces these degrees of freedom.

The good j_0 will be shares in the firm, and the price will be the stock price P_t. I specify that agents believe the stock price to follow the process

$$
\begin{aligned}
\log P_t - \log P_{t-1} &= \mu_t + \eta_t && (1.6)\\
\mu_t &= \mu_{t-1} + \nu_t && (1.7)\\
\text{where } \begin{pmatrix} \eta_t \\ \nu_t \end{pmatrix} &\sim \mathcal{N}\left(-\frac{1}{2}\begin{pmatrix} \sigma_\eta^2 \\ \sigma_\nu^2 \end{pmatrix}, \begin{pmatrix} \sigma_\eta^2 & 0 \\ 0 & \sigma_\nu^2 \end{pmatrix} \right) \; iid,
\end{aligned}
$$

and where all elements of this system except P_t itself are unobserved, as in Adam et al. (2013). When agents use Bayesian updating for their belief $\hat{\mu}_t$ about μ_t, this system can be recast in a pure observable form as follows:

$$
\begin{aligned}
\log P_t - \log P_{t-1} &= \hat{\mu}_{t-1} - \frac{1-g+g^2}{2}\sigma_z^2 + z_t && (1.8)\\
\hat{\mu}_t &= \hat{\mu}_{t-1} - \frac{g^2}{2}\sigma_z^2 + g z_t && (1.9)\\
\text{where } z_t &\sim \mathcal{N}\left(0, \sigma_z^2\right) \; iid.
\end{aligned}
$$

The disturbance z_t is the forecast error on prices. The learning gain g and the variance of the forecast error σ_z^2 are functions of σ_η^2 and σ_ν^2 .

The equations above define the stochastic law of motion of P_t under the subjective belief $\mathcal{P}$. CMCE will completely determine all remaining expectations under this subjective belief; the only choice left is which market j_1 in Definition 1.3 they do not expect to clear. Here, there are two markets besides the stock market: that for regular consumption goods and that for the zero-good. Choosing to relax the expectation of market clearing in either one of these markets will lead to a different belief system and equilibrium stock price process. I will discuss the equilibrium in each case.

1.3.3.1 *Relaxing expectations of consumption goods market clearing*

The first way to introduce CMCE is by relaxing expectations of regular consumption goods market clearing. This way recovers the equilibrium dynamics studied in Adam et al. (2014).

How do conditionally model-consistent expectations determine the beliefs about endogenous model variables? Definition 1.3 requires that the law of motion for exogenous variables **u** be correct, i.e. agents know how to forecast A_t and W_t. Beliefs about everything else have to be consistent with optimal choice (so the household understands $D_t = A_t$) and market clearing in all markets except for the asset market and the market for consumption goods. In particular, agents expect the market for zero-goods to clear, i.e. they expect $C_t^0 = 0$ at all times. The decision rules of agents and the corresponding beliefs have to be determined simultaneously in the following problem:

$$\begin{aligned}
P_t &= \mathbb{E}_t^{\mathcal{P}}\left[\beta\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}(P_{t+1}+A_{t+1})\right] && (1.10)\\
C_t &= W_t + \pi_t && (1.11)\\
\pi_t &= (P_t + A_t)\,S_{t-1} - P_t S_t && (1.12)\\
C_t^0 &= 0 && (1.13)\\
W_t &= \bar{W} && (1.14)\\
\log A_t &= \rho \log A_{t-1} + \varepsilon_t\, \varepsilon_t \overset{\mathcal{P}}{\sim} \mathcal{N}\left(0,\sigma_\varepsilon^2\right) && (1.15)\\
\log P_t &= \log P_{t-1} + \hat{\mu}_{t-1} - \frac{1-g+g^2}{2}\sigma_z^2 + z_t && (1.16)\\
\hat{\mu}_t &= \hat{\mu}_{t-1} - \frac{g^2}{2}\sigma_z^2 + g z_t,\ z_t \overset{\mathcal{P}}{\sim} \mathcal{N}\left(0,\sigma_z^2\right) && (1.17)
\end{aligned}$$

Thus, we recover the same equations as in Adam et al..[5] What is the difference between this subjective model and the rational expectations equilibrium? It is the fact that the market clearing condition $S_t = 1$ is absent. Instead, we are given an (exogenous) subjective law of motion for P_t and operate in a kind of partial equilibrium. The solution is a subjective policy function

$$(S_t, P_t, \hat{\mu}_t, A_t) = h\left(P_{t-1}, \hat{\mu}_{t-1}, S_{t-1}, A_{t-1}, \varepsilon_t, z_t\right)$$

We then solve for the equilibrium by imposing market clearing. This requires the price P_t to be such that $S_t = 1$. Through equation (1.16), there exists a one-to-one mapping between the price P_t and the subjective forecast error z_t. We are effectively looking for a function

5 There is in fact a difference in that here, the wage and dividend process are stationary.

Figure 1.1: Approximation of the equilibrium asset price with learning, 1st approach.

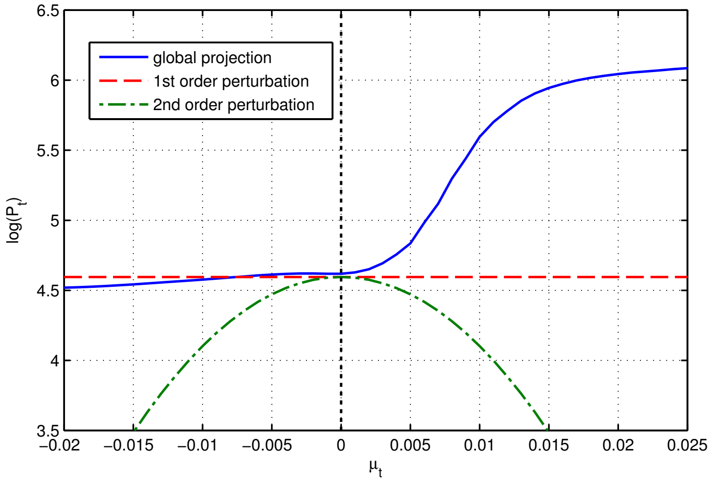

Equilibrium stock price at $A_t = 1$. Black dotted line indicates non-stochastic steady state.

$z_t = r\left(P_{t-1}, A_{t-1}, \hat{\mu}_{t-1}, \varepsilon_t\right)$ for which $S_t = 1$ always. This pins down the evolution of the state variables $\left(P_t, \hat{\mu}_t, A_t\right)$ in equilibrium.

So while agents think that z_t is normally distributed white noise, in reality it is a deterministic function of the state variables and the exogenous shock ε_t. It is in this sense that agents' expectations are not rational expectations. However, conditional on believing in z_t being white noise, agents' expectations are consistent with all equilibrium relations.

In practice, the policy functions can be solved with standard numerical methods. The problem here is that low-order perturbation methods lead to bad approximations. This is problematic because it becomes difficult to embed this type of learning into a larger business cycle model for which the computational costs of global solution methods are very high. Figure 1.1 shows the policy function for the equilibrium price P_t for various values of $\hat{\mu}_t$ at $D_t = 1$ using different solution methods.[6] The blue solid line shows the price function obtained using a global projection method.[7] The equilibrium price is an increasing function of the price growth belief $\hat{\mu}_t$, but the relationship is non-linear. In particular, it is relatively flat around the non-stochastic steady state (where $\hat{\mu}_t = 0$). When the model is linearised around this steady state, the solution is in fact completely independent of the price growth belief (red dashed line). This is clearly a bad approximation, as the goal of introducing learning is to get more

6 The parameters chosen are $\beta = .99$, $\gamma = 2$, $\bar{W} = 20$, $\rho = .95$, $\sigma_\varepsilon = .01$, $\sigma_z = .01$.

7 I use time iteration on the household's Euler equation. The policy function is computed on a grid with 18,360 nodes and linearly interpolated. The expectation is approximated using quadrature methods and 3 Chebyshev nodes.

volatility in prices.[8] Resorting to a second order approximation does not help either. The green dash-dotted line shows the policy function of the equilibrium price approximated to second order around the non-stochastic steady-state. The quadratic coefficient on $\hat{\mu}_t$ is large and negative. This, again, is a bad approximation to the true policy function.

1.3.3.2 *Relaxing expectations of zero-good market clearing*

The second possibility to introduce CMCE is to relax expectations of market clearing in the market for the zero-good. This approach also generates a price which is increasing in the price growth belief and has the advantage that first and second order perturbation methods have better accuracy.

As before, start with the subjective belief (1.8)-(1.9) about P_t; then find expectations and actions that are consistent with optimal choice and market clearing in all markets except that for stock and the zero-good.

The household has to form beliefs about the profits π_t it receives from the investment agency. Since these beliefs must be consistent with clearing of the market for consumption goods, the household must hold beliefs such that he expects his choice to be $C_t = \bar{W} + A_t$ at all times, i.e. $\pi_t = A_t$. The firm still chooses $D_t = A_t$ and all agents must have beliefs consistent with this choice. The investment agency therefore believes that $D_t = A_t$ at all times. For any level of stockholdings, it is always indifferent between paying out profits π_t or transferring zero-goods C_t^0 for the household. But since expectations and optimal choice must be consistent and the household believes $\pi_t = A_t$, the investment agency must choose $\pi_t = A_t$ for Definition 1.3 to hold. It will have the correct belief about household consumption $C_t = A_t + \bar{W}$.

The only variables that remain to be solved for are the investment agency's stock holdings S_t and choice of C_t^0. Under conditionally model-consistent expectations, we can write the investment agency's first order conditions as

$$C_t^0 = (P_t + D_t)\, S_{t-1} - P_t S_t - A_t \tag{1.18}$$

$$S_t \begin{cases} = & 0 \\ \in & [0, \bar{S}] \\ = & \bar{S} \end{cases} \text{if } P_t \gtreqqless \mathbb{E}_t^{\mathcal{P}}\left[Q_{t,t+1}\left(P_{t+1} + A_{t+1}\right)\right] \tag{1.19}$$

8 It can be shown that the independence of the price holds for any process for dividends A_t and non-dividend income W_t, even when they contain a unit root.

Figure 1.2: Approximation of the equilibrium asset price with learning, 2nd approach.

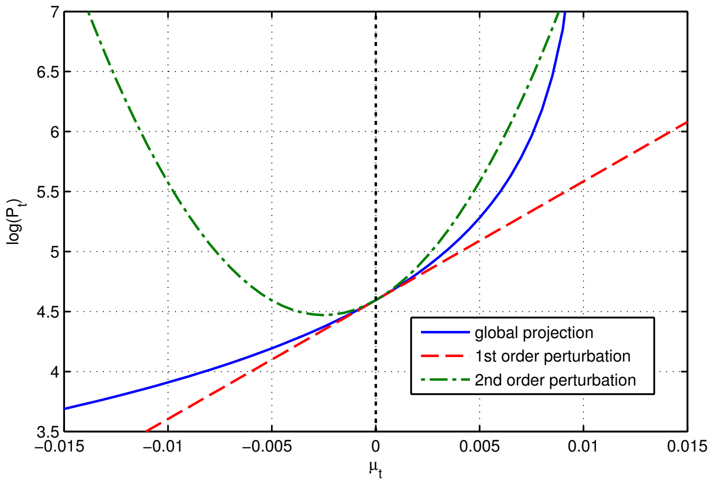

Equilibrium stock price at $A_t = 1$. Black dotted line indicates non-stochastic steady state.

Imposing $S_t = 1$ implies $C_t^0 = 0$, so all markets clear. The equilibrium asset price now has a closed-form solution:

$$\begin{aligned} P_t &= \mathbb{E}_t^{\mathcal{P}}\left[Q_{t,t+1}\left(P_{t+1} + A_{t+1}\right)\right] \\ &= \mathbb{E}_t^{\mathcal{P}}\left[Q_{t,t+1}\left(P_t \exp\left(\hat{\mu}_t - \frac{1-g+g^2}{2}\sigma_z^2 + z_{t+1}\right) + A_{t+1}\right)\right] \\ &= \frac{\beta \mathbb{E}_t\left[\left(\frac{\bar{W}+A_{t+1}}{\bar{W}+A_t}\right)^{-\gamma} A_{t+1}\right]}{1 - \beta \mathbb{E}_t\left[\left(\frac{\bar{W}+A_{t+1}}{\bar{W}+A_t}\right)^{-\gamma}\right] \exp\left(\hat{\mu}_t + \frac{g-g^2}{2}\sigma_z^2\right)}. \end{aligned} \tag{1.20}$$

Using the same perceived law of motion for the price P_t, and a slightly modified version of the model with identical rational expectations solution, we have thus arrived at a different learning equilibrium. The stock price is now globally increasing in the price growth belief $\hat{\mu}_t$. As before, we can compare different approximation methods to this solution. Figure 1.2 shows the values for the equilibrium price P_t at $A_t = 1$ for different values of the price growth belief $\hat{\mu}_t$. The price is increasing in agents' beliefs (the blue dashed line showing a global solution method), but the slope at the non-stochastic steady-state $\hat{\mu}_t = 0$ is now large and positive. A first order linearisation (red dashed line) therefore preserves the general behaviour of the pricing function, and a second order approximation (green dash-dotted line) provides an even better approximation, at least for positive values of $\hat{\mu}_t$.

The combination of the equilibrium pricing functions can be combined with the evolution of beliefs $\hat{\mu}_t$ to examine the asset pricing properties of this simple model. However, the focus of this chapter is on applications of asset price learning to business cycle models. I will therefore move on directly to the study of the first business cycle model. In doing so, I will discuss in turn each of the two approaches developed in this section.

1.4 RBC MODEL

Adam et al. (2013) show that models of stock price learning such as the one described above can replicate several key properties of stock prices, including return volatility and predictability, without the need for complex preferences or high degrees of risk aversion. In this section, I will examine what a learning-based asset pricing theory implies for the real business cycle model.

1.4.1 *Model description*

The model is a standard RBC model and can be thought of as an extension of the model of the last section to an economy with production. The economy is populated by a representative household, a representative firm and an investment agency. There are three physical goods: a regular consumption/investment good that serves as the numéraire, labour which trades in a competitive spot market at the wage w_t, and a "zero-good" in zero net supply. There is no technology to convert zero goods into other goods or vice-versa.

The household likes consumption and dislikes working. As before, it does not value the zero-good. It receives wage payments from selling labour and also receives profits from the investment agency. It maximises utility as follows:

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

$$\text{s.t. } C_t = w_t L_t + \pi_t.$$

Its first order condition is:

$$w_t = \eta C_t^{\gamma} L_t^{\phi}. \tag{1.21}$$

The investment agency solves the following problem:

$$\max_{\left(S_t,\pi_t,C_t^0\right)_{t=0}^{\infty}} \mathbb{E}^{\mathcal{P}} \sum_{t=0}^{\infty} Q_{0,t} \left(\pi_t + C_t^0\right)$$

$$\begin{aligned} \text{s.t. } \pi_t + C_t^0 + P_t S_t &\leq (P_t + D_t)\, S_{t-1} \\ S_t &\in \left[0, \bar{S}\right]. \end{aligned}$$

where the household discount factor is defined as $Q_{t,s} = \beta^{s-t}\left(C_s/C_t\right)^{-\gamma}$ for $s, t \geq 0$. Again, the investment agency has the option to either pay out profits π_t to the household or transfer zero-goods C_t^0 to it. In equilibrium it has to be $C_t^0 = 0$ but the investment agency may not be aware of this. The first order conditions of the investment agency are the same as in the last section.

The representative firm owns the capital stock K_{t-1} at the beginning of the period. It combines this capital with labour L_t to produce output using a Cobb-Douglas production function. In the process of production, the capital stock depreciates at the rate δ. The firm then decides on next period's capital stock K_t. Any remaining earnings are paid out as dividends to shareholders. The firm maximises the discounted sum of its dividend payments, discounted at the household rate:

$$\max_{(K_t, L_t, D_t)_{t=0}^{\infty}} \mathbb{E}^{\mathcal{P}} \sum_{t=0}^{\infty} Q_{0,t} D_t$$

$$\begin{aligned} \text{s.t. } D_t &= Y_t - w_t L_t - K_t + (1-\delta) K_{t-1} \\ Y_t &= A_t K_t^{\alpha} L_t^{1-\alpha} \end{aligned}$$

where total factor productivity A_t is a stochastic exogenous process with the law of motion

$$\log A_t = (1-\rho)\bar{A} + \rho A_{t-1} + \varepsilon_t. \tag{1.22}$$

and where $\varepsilon_t \sim \mathcal{N}\left(0, \sigma_\varepsilon^2\right)$ is an iid white noise process. The first order conditions of the firm are:

$$1 = \mathbb{E}_t^{\mathcal{P}}\left[Q_{t,t+1}\left(\alpha A_{t+1}\left(\frac{L_{t+1}}{K_t}\right)^{1-\alpha} + 1 - \delta\right)\right] \tag{1.23}$$

$$w_t = (1-\alpha) A_t \left(\frac{K_{t-1}}{L_t}\right)^{\alpha} \tag{1.24}$$

The model description is completed with the market clearing conditions. The number of shares owned by the household has to be $S_t = 1$ (no share issuance or buybacks). Total output has to equal the sum of consumption and investment: $Y_t = C_t + K_t - (1-\delta) K_{t-1}$.

1.4.2 *Learning equilibrium*

I will now add learning about stock prices P_t to the model with conditionally model-consistent expectations. As in the previous section, agents believe the price P_t to evolve according to the following law of motion:

$$\log P_t - \log P_{t-1} = \hat{\mu}_{t-1} - \frac{1-g+g^2}{2}\sigma_z^2 + z_t \tag{1.25}$$

$$\hat{\mu}_t = \hat{\mu}_{t-1} - \frac{g^2}{2}\sigma_z^2 + g z_t \tag{1.26}$$

$$\text{where } z_t \sim \mathcal{N}\left(0, \sigma_z^2\right) \; iid.$$

I then impose conditionally model-consistent expectations with respect to P_t, which are consistent with clearing of all markets except the market for stocks and one other market. Again, choosing a different market which agents do not expect to clear leads to different equilibrium outcomes. I will discuss both approaches from the last section in turn: relaxing expectations of consumption goods market clearing, and relaxing expectations of zero-good market clearing.

1.4.2.1 *Relaxing expectations of consumption goods market clearing*

Given the law of motion of productivity and the subjective beliefs about stock prices, I need to solve for a complete system of expectations that are consistent with optimal choice and market clearing in all markets except that for stocks and the regular consumption good. This amounts to solving a recursive model comprising the following equations:

$$
\begin{aligned}
P_t &= \mathbb{E}_t^{\mathcal{P}}\left[\beta\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}(P_{t+1}+D_{t+1})\right] && (1.27)\\
C_t &= w_t L_t + \pi_t && (1.28)\\
w_t &= \eta C_t^{\gamma} L_t^{\phi} && (1.29)\\
\pi_t &= (P_t + D_t)\, S_{t-1} - P_t S_t && (1.30)\\
C_t^0 &= 0 && (1.31)\\
w_t &= (1-\alpha)\, A_t \left(\frac{K_{t-1}}{L_t}\right)^{\alpha} && (1.32)\\
D_t &= \alpha A_t K_{t-1}^{\alpha} L_t^{1-\alpha} - w_t L_t + (1-\delta)\, K_{t-1} - K_t && (1.33)\\
1 &= \mathbb{E}_t^{\mathcal{P}}\left[\beta\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}\left(\alpha A_{t+1}\left(\frac{L_{t+1}}{K_t}\right)^{1-\alpha} + 1 - \delta\right)\right] && (1.34)\\
\log A_t &= (1-\rho)\, \bar{A} + \rho \log A_{t-1} + \varepsilon_t,\ \varepsilon_t \overset{\mathcal{P}}{\sim} \mathcal{N}\left(0, \sigma_\varepsilon^2\right) && (1.35)\\
\log P_t &= \log P_{t-1} + \hat{\mu}_{t-1} - \frac{1-g+g^2}{2}\sigma_z^2 + z_t && (1.36)\\
\hat{\mu}_t &= \hat{\mu}_{t-1} - \frac{g^2}{2}\sigma_z^2 + g z_t,\ z_t \overset{\mathcal{P}}{\sim} \mathcal{N}\left(0, \sigma_z^2\right) && (1.37)
\end{aligned}
$$

In this formulation, the only equilibrium condition that is missing is the market clearing condition for regular consumption/investment goods. It is instead replaced with the exogenous price process for P_t. Where the rational expectations solution has two state variables (A_{t-1} and K_{t-1}) and one shock (ε_t), the perceived policy function has five state variables (A_{t-1}, K_{t-1}, S_{t-1}, P_{t-1} and $\hat{\mu}_t$) and two shocks (ε_t and z_t). The subjective policy function takes the form

$$(A_t, K_t, S_t, P_t, \hat{\mu}_t) = h\left(A_{t-1}, K_{t-1}, S_{t-1}, P_{t-1}, \hat{\mu}_{t-1}, \varepsilon_t, z_t\right).$$

As before, I can solve for the equilibrium by imposing market clearing $S_t = 1$ (the goods market then automatically clears by Walras' law).

Table 1.1: Calibration of the RBC model.

	parameter	value		parameter	value
α	capital factor share	0.33	γ	intertemporal elasticity of substitution	1
δ	depreciation rate	0.025	ϕ	inverse Frisch elasticity	0.33
ρ	persistence of productivity shock	0.95	η	utility weight on leisure	0.876
β	discount factor	0.99	g	learning gain	0.009
σ_ε	std. deviation of productivity shock	0.007	σ_z	perceived std. deviation of stock prices	0.001

That is, I am looking for a function $z_t = r\left(A_{t-1}, K_{t-1}, S_{t-1}, P_{t-1}, \hat{\mu}_{t-1}, \varepsilon_t\right)$ for which $S_t = 1$ always holds. This then leads to the actual policy function.

It has been demonstrated in the last section that perturbation methods (at least of a low order) are poor approximations to the model solution under learning. This is of course in contrast to the rational expectations version of the RBC model for which a first order approximation is already very accurate. Here, I have to solve the model using global projection methods.[9]

1.4.2.2 *Calibration and results*

Table 1.1 summarises the calibration of the model, which follows the standard RBC calibration. The only parameters specific to learning are the learning gain g and the perceived volatility of stock prices σ_z. I chose a combination for those two parameters for which the volatility of stock returns is broadly similar to the data.

Figure 1.3 shows impulse responses after a one-standard deviation technology shock (increase of A_0 by 0.7%). The red solid line is the impulse response under learning while the blue-dashed line is the response of the rational expectations version of the model. Output, investment, consumption and employment all rise after the shock. However, at impact, consumption initially rises by less under learning

9 Again, I use time iteration on the household's Euler equation and the firm's investment equation. The policy function h is computed on a grid with 35,937 nodes and linearly interpolated. The expectation is approximated using quadrature methods and 3 Chebyshev nodes.

Figure 1.3: Impulse responses to a productivity shock in the RBC model.

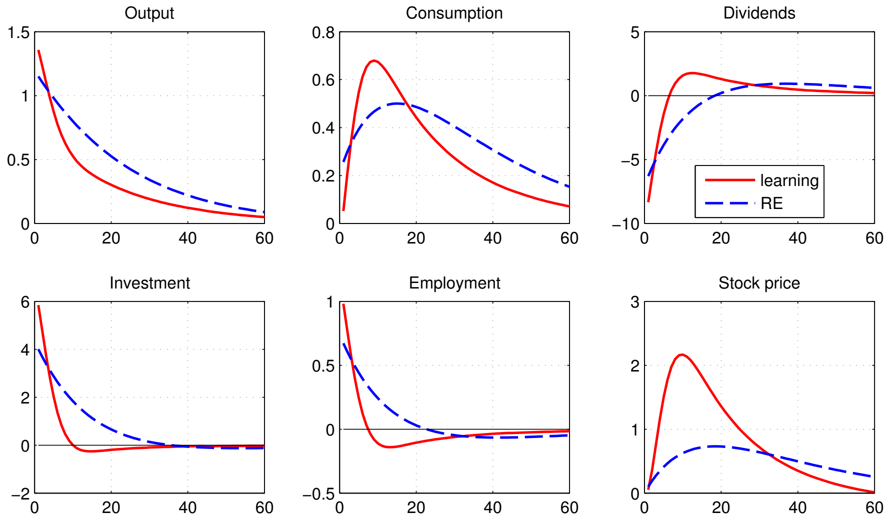

Impulse responses to a one-standard deviation innovation in ε_t. Output Y_t, consumption C_t, dividends D_t, investment I_t, employment L_t, and the stock price P_t are in 100*log deviations.

than under rational expectations. Since this expands labour supply, employment and output are higher, allowing for a larger rise in investment. Because of larger costs of investment, dividends fall by more than under rational expectations. Over time, however, agents' price growth beliefs $\hat{\mu}_t$ increase and the equilibrium stock price rises. This rise (and expectations of further rising prices) acts as a positive wealth effect on households. They increase consumption and decrease labour supply at the expense of output and investment by firms. Intuitively, households expect higher returns from stocks. Firms can deliver higher marginal returns by cutting down on investment, increasing the marginal product of capital. Lower investment increases dividend payments, allowing households to consume more and reduce their labour supply. They are happy for firms to invest less because they expect to expand their consumption in the future through their asset income (failing to realise that in general equilibrium, the price of those assets is determined by a fixed supply). Therefore, the response of output, investment and employment fall below their rational expectations counterpart, while the opposite holds for consumption.

An increase in the stock price P_t and/or the price growth belief $\hat{\mu}_t$ acts like a "news shock", essentially an expectation of future higher income. It is well known (Beaudry and Portier, 2007) that the real business cycle model cannot produce comovement of output, consumption and employment in response to news shocks. The learning model also suffers from this comovement problem.

Table 2.4 compares second moments across quarterly U.S. data, the rational expectations version of the model (Column 1) and the learning version, for two values of the learning gain (Column 2). The learning model using the second approach (relaxing expectations of market clearing for zero-goods, Column 3) will be described further below.

The rational expectations RBC model is known to match relative volatilities of business cycle aggregates and their correlation with output well. But asset prices are too smooth relative to the data: the standard deviation of stock returns in the data is 33.5%, but less than a fraction of a percent in the RE model. This is despite the fact that dividends are twice as volatile as in the data. They are also countercyclical: When firms in the RBC model increase investment after a positive productivity shock, this directly reduces the dividends they pay. Variation in the price-dividend ratio therefore comes almost entirely from movements in dividends. The RBC model also does not capture return predictability adequately. At the four quarter horizon, the P/D-ratio positively predicts future returns, at odds with the data. At the 20-quarter horizon, the prediction reverses. But this predictability comes mostly from the mean reversion in dividends.

The model with learning improves the asset price moments, but at the expense of a worse fit on business cycle moments. The volatility of returns and the P/D-ratio is much higher, and predictability at the 4-quarter horizon is in line with the data (although at the 20-quarter horizon it disappears). The problem is that consumption, investment and employment are now excessively volatile relative to output. What's more, consumption is negatively correlated with output. This is the comovement problem described earlier on.

Overall, the model properties are not very appealing. Could a slight modification of the model solve the comovement issue and at the same time preserve the asset price properties? My tentative answer is no. The literature on news shocks has pointed out that frictionless neoclassical models can only generate positive comovement in response to expectational shifts with rather unusual modifications. An exception is Jaimovich and Rebelo (2009): They use a combination of adjustment costs, variable capacity utilisation and Greenwood-Hercowitz-Huffman preferences. But in their model, positive comovement rests on firms which expect higher productivity in the future to frontload investment in order to avoid adjustment costs. With learning about stock prices however, a change in expectations about stock prices does not do anything to firms' expectations about productivity. Therefore, it is not to be expected that the comovement problem can be solved here without introducing additional frictions. This leads me to consider models with search frictions in the labour market in the next section.

Table 1.2: Moments in the data and across RBC model specifications.

	moment	data	(1) RE	(2) learn I	(3) learn II
output volatility	$\sigma_{hp}(Y_t)$	1.43%	1.50%	3.85%	1.50%
relative volatility	$\sigma_{hp}(C_t)/\sigma_{hp}(Y_t)$	0.60	0.27	0.97	0.27
	$\sigma_{hp}(I_t)/\sigma_{hp}(Y_t)$	2.90	3.54	4.76	3.54
	$\sigma_{hp}(D_t)/\sigma_{hp}(Y_t)$	2.99	5.53	5.39	5.53
cross-correlation	$\rho_{hp}(C_t, Y_t)$	0.94	0.84	-0.82	0.84
	$\rho_{hp}(I_t, Y_t)$	0.95	0.99	0.87	0.99
asset price volatility	$\sigma_{hp}\left(\frac{P_t}{D_t}\right)$	11.4%	8.56%	14.2%	7.72%
	$\sigma\left(\frac{P_t}{D_t}\right)$	40.8%	14.1%	22.0%	14.1%
	$\sigma\left(R^{stock}_{t,t+1}\right)$	33.5%	0.38%	25.3%	42.6%
return predictability	$\rho\left(\frac{P_t}{D_t}, R^{stock}_{t,t+4}\right)$	-0.22	0.38	-0.19	-0.62
	$\rho\left(\frac{P_t}{D_t}, R^{stock}_{t,t+20}\right)$	-0.44	-0.26	0.09	-0.66

Quarterly US data 1962Q1-2012Q4. Consumption C_t consists of 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. The stock price index P_t is the S&P500. All variables deflated by the GDP deflator. Returns R_t^{stock} are annualised real quarterly stock returns including dividends. The subscript "hp" indicates application of the Hodrick-Prescott filter with smoothing parameter 1600.

1.4.2.3 *Relaxing expectations of zero-good market clearing*

I will now discuss the solution of the model using the second approach. Here, expectations are required to be consistent with market clearing for regular consumption/investment goods, but instead do not need to be consistent with clearing of the market for the zero-good. The system of equations governing the subjective expectations and policy functions then reads as follows:

$$C_t = A_t K_{t-1}^{\alpha} L_t^{1-\alpha} \tag{1.38}$$

$$w_t = \eta C_t^{\gamma} L_t^{\phi} \tag{1.39}$$

$$w_t = (1-\alpha) A_t \left(\frac{K_{t-1}}{L_t}\right)^{\alpha} \tag{1.40}$$

$$D_t = \alpha A_t K_{t-1}^{\alpha} L_t^{1-\alpha} - w_t L_t + (1-\delta) K_{t-1} - K_t \tag{1.41}$$

$$1 = \mathbb{E}_t^{\mathcal{P}} \left[\beta \left(\frac{C_{t+1}}{C_t}\right)^{-\gamma} \left(\alpha A_{t+1} \left(\frac{L_{t+1}}{K_t}\right)^{1-\alpha} + 1 - \delta\right)\right] \tag{1.42}$$

$$\log A_t = (1-\rho) \bar{A} + \rho \log A_{t-1} + \varepsilon_t,\ \varepsilon_t \sim \mathcal{N}\left(0, \sigma_\varepsilon^2\right) \tag{1.43}$$

$$\log P_t = \log P_{t-1} + \hat{\mu}_{t-1} - \frac{1-g+g^2}{2}\sigma_z^2 + z_t \tag{1.44}$$

$$\hat{\mu}_t = \hat{\mu}_{t-1} - \frac{g^2}{2}\sigma_z^2 + g z_t,\ z_t \sim \mathcal{N}\left(0, \sigma_z^2\right) \tag{1.45}$$

It turns out that the allocations produced by this approach to learning are identical to those under rational expectations. The first four equations together with the correct law of motion for A_t define exactly the same choices for investment, employment and consumption as under rational expectations. The realisation of the stock price P_t appears in none of the optimality conditions. This is because CMCE requires consistency of expectations with optimal choices and market clearing, and since I have kept market clearing for consumption/investment goods, agents actually hold rational expectations with respect to real variables. They do not hold rational expectations only with respect to stock prices. The investment agency also expects to buy non-zero quantities of C_t^0 when prices P_t do not follow the rational expectations path.

In equilibrium $C_t^0 = 0$ has to hold every period. This amounts to imposing the Euler equation of the investment agency:

$$P_t = \mathbb{E}_t^{\mathcal{P}} \left[\beta \left(\frac{C_{t+1}}{C_t}\right)^{-\gamma} (P_{t+1} + D_{t+1})\right] \tag{1.46}$$

$$= \frac{\beta \mathbb{E}_t^{\mathcal{P}} \left[\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma} D_{t+1}\right]}{1 - \beta \mathbb{E}_t^{\mathcal{P}} \left[\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}\right] \exp\left(\hat{\mu}_t + \frac{g-g^2}{2}\sigma_z^2\right)} \tag{1.47}$$

and solving this for the equilibrium forecast error z_t.

Figure 1.4: Dividend and price response to a productivity shock in the RBC model.

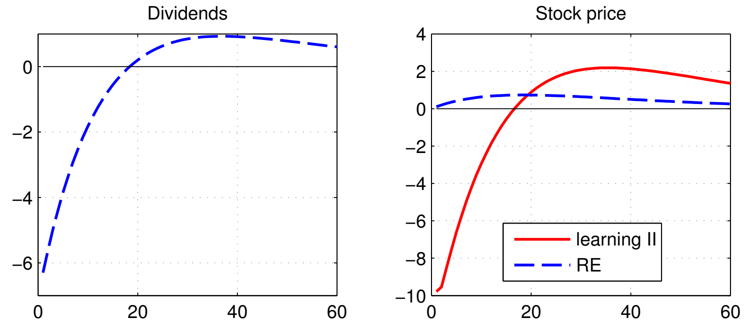

Impulse responses to a one-standard deviation innovation in ε_t. Dividends D_t and the stock price P_t are in 100*log deviations.

The impact of stock price learning in this version of the model is therefore entirely one-sided: While allocations affect stock prices, stock prices do not feed back into allocations. Figure 1.4 plots impulse response functions for dividends and stock prices. Dividends fall initially after a positive technology shock, but under rational expectations, the stock price still rises slightly in anticipation of higher dividends further in the future. By contrast, with this version of learning, the stock price reacts much more and falls substantially.

The last column of Table 2.4 shows the second moments of this second version of the learning model. By construction, the business cycle moments are identical to the rational expectations version. Return volatility and predictability are much closer to the data. However, stock prices are countercyclical in this version of the model.

So while the good business cycle properties of the RBC model are preserved with this way of learning, and stock prices are volatile as in the data, this still doesn't seem to be a good description of the joint behaviour of stock prices and business cycles.

1.5 SEARCH MODEL

From the previous discussion, it should have become clear that the RBC model augmented with learning suffers from counterfactual comovement of key variables, and that the comovement problems are related to similar difficulties encountered in the literature on news shocks. I now turn to the search and matching model of unemployment. This model has become a standard tool to study business cycles with more realistic labour markets. More importantly, it is known to generate the right comovement in response to news shocks (den Haan and Kaltenbrunner, 2009).

It will turn out that positive comovement is possible in response to fluctuations in beliefs of learning investors as well. Adding learning improves the asset price properties of the model and amplifies productivity shocks with positive comovement. What's more, learning will also address the "Shimer puzzle": The observation by Shimer (2005) that the search model is not able to generate sufficient volatility in unemployment unless substantial wage rigidity is assumed. I show that the degree of wage rigidity necessary to get the right volatility of unemployment decreases substantially when one allows for learning in the stock market.

1.5.1 *Model description*

The model is a standard search and matching model in discrete time. The economy is populated by workers, entrepreneurs, and an investment agency. All agents can perfectly insure idiosyncratic risk, as at the end of each period they pool their net revenue within a representative household. There are three physical goods: labour, a consumption/investment good which serves as the numéraire, and a zero-good as in the previous section. Additionally, there is a rental market for capital and a market for shares in operating firms. All markets are competitive except for the labour market, which is subject to search frictions.

1.5.1.1 *Firms and entrepreneurs*

At each point in time, there are n_t firms operating in the economy. Each firm employs one worker, so n_t is also the level of employment. A firm rents k_t units of capital in a competitive market at the price r_t. Its production function is

$$y_t = a_t^{1-\alpha} k_t^{\alpha} \tag{1.48}$$

$$\log a_t = \rho a_{t-1} + (1-\rho)\,\bar{a} + \varepsilon_t \tag{1.49}$$

where a_t is the level of productivity common to all firms, which evolves according to an AR(1) process in logarithms with iid innovations $\varepsilon_t \sim \mathcal{N}\left(0, \sigma^2\right)$. The firm and the worker share the surplus from production:

$$\Omega_t = \max_{k_t} y_t - r_t k_t = (1-\alpha)\, y_t \tag{1.50}$$

according to the formula in den Haan and Kaltenbrunner (2009):

$$w_t = (1-\eta)\left((1-\omega)\,\Omega_t + \omega\bar{\Omega}\right). \tag{1.51}$$

where $\bar{\Omega} = \mathbb{E}^{\mathcal{P}}\left[\Omega_t\right]$ is the unconditional expectation of the surplus. The parameter ω governs the amount of wage stickiness. When $\omega = 1$, the wage is completely acyclical, while it comoves perfectly with labour productivity when $\omega = 0$. The parameter η governs the share of the surplus that accrues to firms. The firm pays it out as dividends to shareholders:

$$\begin{aligned} d_t &= \Omega_t - w_t \\ &= \eta \Omega_t + \omega (1 - \eta) (\Omega_t - \bar{\Omega}) \end{aligned} \tag{1.52}$$

Firms and workers separate at the exogenous rate s. There is no endogenous job destruction.

Each period, entrepreneurs decide how many firms to create. Creating a firm entails finding a worker by posting a vacancy, which entails paying a cost κ. With probability q_t, a worker is matched and the firm immediately starts production. Ownership in the firm can be sold in the form of a unit mass of shares in the stock market.[10] The market value of a firm is denoted p_t. Free entry implies that the cost and expected gain from starting a project have to be equal:

$$\kappa = q_t p_t \tag{1.53}$$

1.5.1.2 *Household*

The representative household consists of a unit mass of workers who always participate in the labour market. It takes the aggregate employment level n_t as given. It chooses consumption C_t and capital K_t to solve the following maximisation problem:

$$\max_{C_t, K_t} \mathbb{E}^{\mathcal{P}} \sum_{t=0}^{\infty} \frac{C_t^{1-\gamma}}{1-\gamma}$$

$$\text{s.t. } C_t + K_t = n_t w_t + (1 - \delta) K_{t-1} + r_t K_{t-1} + \pi_t$$

The capital stock of the last period is rented out to firms at the rental rate r_t and depreciates at the rate δ. Additionally, the household receives payments π_t from the investment agency which it owns. The first order condition for the household problem is

$$1 = \mathbb{E}_t^{\mathcal{P}} \left[Q_{t,t+1} (r_{t+1} + 1 - \delta) \right] \tag{1.54}$$

where the household discount factor is defined as $Q_{t,s} = \beta^{s-t} (C_s / C_t)^{-\gamma}$.

1.5.1.3 *Investment agency*

The investment agency trades shares of firms in the stock market. At the beginning of each period, it owns a number S_{t-1} of shares in firms. A fraction s of firms exit at random and thus their shares become worthless. The remaining shares each trade at the price p_t. The agency can then buy a number S_t of new shares at the price p_t and immediately receives dividend payments D_t on those shares.

10 It is inconsequential whether entrepreneurs sell off ownership in operating firms to the household or keep the shares, since all income is pooled at the end of a period.

Thus, p_t is the cum-dividend share price. The investment agency solves the following maximisation problem:

$$\max_{S_t, \pi_t, x_t} \mathbb{E}^{\mathcal{P}} \sum_{t=0}^{\infty} Q_{0,t} \left(\pi_t + C_t^0 \right)$$

$$\begin{aligned} \text{s.t. } \pi_t + C_t^0 + p_t S_t &= (1-s)\, p_t S_{t-1} + d_t S_t \\ S_t &\in [0, \bar{S}] \end{aligned}$$

Here, π_t are the profits it pays to the household. Profits are discounted at the household rate $Q_{t,t+1}$. In addition, the agency also values transferring C_t^0 of the zero-good to the household. Paying profits to the household and transferring the zero-good are perfectly substitutable activities in the optimisation problem. Holdings of shares are limited from below by zero (no short-selling) and from above by some constant $\bar{S} > 1$. As before, this constraint will not be binding in equilibrium but guarantees the existence of an equilibrium for arbitrary beliefs about the stock price p_t.

The first order condition of the investment agency is

$$S_t = 0 \quad \text{if} \quad p_t + d_t > (1-s)\, \mathbb{E}_t^{\mathcal{P}} \left[Q_{t,t+1} p_{t+1} \right], \tag{1.55}$$

$$S_t \in [0, \bar{S}] \quad \text{if} \quad p_t + d_t = (1-s)\, \mathbb{E}_t^{\mathcal{P}} \left[Q_{t,t+1} p_{t+1} \right], \tag{1.56}$$

$$S_t = \bar{S} \quad \text{if} \quad p_t + d_t < (1-s)\, \mathbb{E}_t^{\mathcal{P}} \left[Q_{t,t+1} p_{t+1} \right]. \tag{1.57}$$

1.5.1.4 *Matching and market clearing*

Each period, the number of matches between vacancies and workers is given by a Cobb-Douglas matching function:

$$m_t = \nu u_t^{\mu} v_t^{1-\mu} \tag{1.58}$$

where v_t is the number of vacancies and u_t is the number of unemployed workers. Consequently, the probability that an unemployed worker becomes employed is given by $f_t = m_t / u_t$, and the probability that a vacancy is filled with a worker is given by $q_t = m_t / v_t$. Aggregate employment consists in workers who were employed last period, minus separations, plus newly matched workers:

$$n_t = (1-s)\, n_{t-1} + f_t u_t \tag{1.59}$$

$$u_t = 1 - n_{t-1} \tag{1.60}$$

The remaining markets are competitive. The zero-good, as the name says, is in zero net supply. The number of shares owned by the investment agency has to equal the number of firms in operation. The aggregate demand for capital by firms has to equal the capital stock owned by the household. Finally, total output has to equal the sum of consumption, investment and vacancy creation costs. In sum, the market clearing conditions of the model are given by:

$$x_t = 0 \tag{1.61}$$
$$S_t = n_t \tag{1.62}$$
$$K_{t-1} = k_t n_t \tag{1.63}$$
$$y_t n_t = C_t + K_t - (1-\delta) K_{t-1} + \kappa v_t \tag{1.64}$$

1.5.2 *Learning equilibrium*

I now add learning about firm value p_t into the model. I directly use the second approach of the previous sections, requiring conditionally model-consistent expectations in which agents do not expect market clearing for stocks and the zero-good. In contrast to the RBC model, learning done in this way will have an effect on allocations, since variations in stock prices influence the incentives of entrepreneurs to create firms.

I impose the same perceived law of motion for the stock price p_t and solve for an internally rational equilibrium with conditionally model-consistent expectations, removing the market clearing conditions for the stock market and the zero-good. The system of equations governing the subjective expectations reads as follows:

$$y_t n_t = C_t + K_t - (1-\delta) K_{t-1} + \kappa_v \theta_t u_t \tag{1.65}$$
$$y_t = a_t k_t^{\alpha} \tag{1.66}$$
$$r_t k_t = \alpha y_t \tag{1.67}$$
$$\Omega_t = (1-\alpha) y_t \tag{1.68}$$
$$w_t = (1-\eta)\left(\omega \Omega_t + (1-\omega) \bar{\Omega}\right) \tag{1.69}$$
$$d_t = \Omega_t - w_t \tag{1.70}$$
$$\kappa_v = q_t P_t \tag{1.71}$$
$$K_{t-1} = k_t n_t \tag{1.72}$$
$$1 = \mathbb{E}_t^{\mathcal{P}}\left[\beta\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}(r_{t+1} + 1 - \delta)\right] \tag{1.73}$$
$$n_t = (1-s) n_{t-1} + f_t u_t \tag{1.74}$$
$$u_t = 1 - n_{t-1} \tag{1.75}$$
$$q_t = \kappa_m \theta^{-\mu} \tag{1.76}$$
$$f_t = \theta_t q_t \tag{1.77}$$

$$\log a_t = (1-\rho)\bar{a} + \rho \log a_{t-1} + \varepsilon_t,\ \varepsilon_t \overset{\mathcal{P}}{\sim} \mathcal{N}\left(0, \sigma_\varepsilon^2\right) \tag{1.78}$$
$$\log p_t = \log p_{t-1} + \hat{\mu}_{t-1} - \frac{1-g+g^2}{2}\sigma_z^2 + z_t \tag{1.79}$$
$$\hat{\mu}_t = \hat{\mu}_{t-1} - \frac{g^2}{2}\sigma_z^2 + g z_t,\ z_t \overset{\mathcal{P}}{\sim} \mathcal{N}\left(0, \sigma_z^2\right) \tag{1.80}$$

In this formulation, the only equilibrium condition that is missing is the market clearing condition for the zero good and the stock

Table 1.3: Calibrated parameters of the search model.

	parameter	value		parameter	value
α	capital factor share	0.33	μ	matching fn. elasticity	0.5
δ	depreciation rate	0.0087	s	separation rate	0.0274
$\bar{A}$	steady state productivity level	0.0379	κ_m	scaling of matching fn.	1
ρ	persistence of productivity shock	0.98	κ_v	vacancy costs	0.042
β	discount factor	0.997	η	entrepreneur surplus share	0.0255
γ	intertemporal elasticity of substitution	1.5	σ_z	perceived std. deviation of stock prices	0.001

market (where $S_t = n_t$ must hold in equilibrium). This is instead replaced with the exogenous price process for p_t. This price has an impact on allocations through the job creation condition. Where the rational expectations solution has three state variables (a_{t-1}, K_{t-1} and n_{t-1}) and one shock (ε_t), the perceived policy function has five state variables (a_{t-1}, K_{t-1}, n_t, p_{t-1} and $\hat{\mu}_t$) and two shocks (ε_t and z_t). The subjective policy function takes the form

$$(a_t, K_t, n_t, p_t, \hat{\mu}_t) = h\left(a_{t-1}, K_{t-1}, n_{t-1}, p_{t-1}, \hat{\mu}_{t-1}, \varepsilon_t, z_t\right).$$

I can solve for the equilibrium by imposing market clearing in the stock market. The investment agency will hold a quantity of stocks $S_t = n_t$ if its Euler equation holds with equality:

$$p_t + d_t = (1-s)\,\mathbb{E}_t^{\mathcal{P}}\left[\beta\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma} p_{t+1}\right]$$

That is, I solve for a function $z_t = r\left(a_{t-1}, K_{t-1}, n_{t-1}, p_{t-1}, \hat{\mu}_{t-1}, \varepsilon_t\right)$ for which the above equation always holds when substituting in the policy function h and the equilibrium relations above.

1.5.3 *Choice of parameters*

The first set of parameters is set to standard values, summarised in Table 1.3. The model is cast at monthly frequency. The factor share of capital is set to $\alpha = 0.33$. The depreciation rate $\delta = 0.0087$ corresponds to an annual depreciation rate of 10 percent. The steady state productivity level $\bar{A}$ is a scaling parameter and is set such that the steady state capital level equals one. The persistence of the productivity shock is set to $\rho = 0.98$. When the productivity process is aggregated up to quarterly frequency, its first autocorrelation is then equal to the standard value 0.95. The discount factor $\beta = 0.997$ implies an annual steady state discount rate of 4 percent. The intertemporal elasticity of substitution is set to $\gamma = 1.5$. This value mainly determines how a persistent increase in output is split between investment and consumption over time. For the chosen value, consumption and investment both rise after a positive change in stock price expectations.

The matching elasticity is set to $\mu = 0.5$, taken from Petrongolo and Pissarides (2001). The parameter κ_m is a scaling parameter and I normalise it to one. The separation rate s and the vacancy cost κ_v are then chosen such that the monthly job finding rate equals 45.4% as estimated by Shimer (2005) and the average unemployment rate is 5.7%. The entrepreneurial share of labour productivity is set to $\eta = 0.0255$, which is the value estimated by Hagedorn and Manovskii (2008). This value is far from the Hosios (1990) efficiency condition: the entrepreneur's share is far smaller than the elasticity of the matching function with respect to vacancies (which is one half). It implies that average job creation is inefficiently low, and this is what allows the model to exhibit positive comovement of output, investment and consumption in response to expectational shocks. Finally, the perceived variance of stock price shocks is set to $\sigma_z = 0.001$.

The remaining parameters are the standard deviation of the productivity shock σ_a, the degree of wage stickiness ω and the learning gain g. I set these parameters to directly match certain second moments in U.S. data. The parameters are estimated by simulated method of moments to match the standard deviation of output and employment; and additionally under learning the standard deviation of the price-dividend ratio. Table 1.4 presents the estimation results. The standard deviation of the productivity shock is estimated at the same magnitude both under learning and rational expectations and corresponds to a value at quarterly aggregation of 0.008. The estimated value for the wage rigidity parameter is $\omega = 0.134$ under learning, but $\omega = 0.592$ under rational expectations. Under learning, the gain g is estimated at $g = 0.0244$ which is in line with results in Adam et al. (2014).

1.5.4 *Results*

1.5.4.1 *Moments*

Table 1.5 compares second moments of U.S. data with the model under learning and under rational expectations. By construction, the

Table 1.4: Estimated parameters of the search model.

parameter	σ_a	ω	g
learning	0.0048	0.134	0.0244
	$(5.91 \cdot 10^{-4})$	(0.040)	(0.0021)
RE	0.0048	0.592	-
	$(6.07 \cdot 10^{-4})$	(0.100)	-

Parameters as estimated by simulated method of moments. Asymptotic standard errors in parentheses. Targets of the moment matching are the standard deviation of the HP-filtered output and unemployment series as described in Table 1.5; and additionally under learning, the standard deviation of the HP-filtered price-dividend ratio.

model matches the volatility of detrended output and unemployment and, additionally under learning, the detrended price/dividend-ratio. Under learning however, much less wage rigidity is needed to achieve the volatility of unemployment in the data. The larger wage rigidity under rational expectations translates directly into a lower volatility of wages and a higher volatility of dividends (see the panel "relative volatility" in the table). The standard deviations as well as the correlations with output of consumption and investment are in line with the data in both cases.

The asset price moments in the bottom half of Table 1.5 reveal large differences between the learning and the rational expectations version of the model. In both both cases, stock returns are actually more volatile than in the data, even though the price dividend ratio itself is less volatile. This indicates that, although prices move a lot, dividends move together with prices more than in the data, stabilising the price dividend ratio. This is the case both under learning and under rational expectations. Note that the rational expectations version generates more than enough return volatility on its own, but that this is achieved through counterfactually large dividend volatility. Turning to return predictability, the learning model does better than rational expectations: while predictability by the P/D ratio is not as strong as in the data, it has at least the right sign. Under rational expectations, high stock prices tend to be followed by high returns, at odds with the data.

1.5.4.2 *Impulse response functions*

Figure 1.5 plots the impulse responses of the model economy after a positive productivity shock. The shock has a size of one standard deviation (i.e. productivity a_t increases by 0.48% on impact). The red

Table 1.5: Moments in the data and across search model specifications.

	moment	data	(1) learning	(2) RE
output volatility	$\sigma_{hp}(Y_t)$	1.52%	1.52%†	1.52%†
unemployment volatility	$\sigma_{hp}(u_t)$	0.77%	0.77%†	0.77%†
relative volatility	$\sigma_{hp}(C_t)/\sigma_{hp}(Y_t)$	0.55	0.39	0.30
	$\sigma_{hp}(I_t)/\sigma_{hp}(Y_t)$	2.70	2.40	2.66
	$\sigma_{hp}(w_t)/\sigma_{hp}(Y_t)$	0.59	0.46	0.20
	$\sigma_{hp}(D_t)/\sigma_{hp}(Y_t)$	4.15	3.57	12.71
correlation with output	$\rho_{hp}(u_t, Y_t)$	-0.88	-0.94	-0.99
	$\rho_{hp}(C_t, Y_t)$	0.85	0.94	0.95
	$\rho_{hp}(I_t, Y_t)$	0.91	0.99	0.99
asset price volatility	$\sigma_{hp}\left(\frac{P_t}{D_t}\right)$	11.2%	11.2%†	4.19%
	$\sigma\left(\frac{P_t}{D_t}\right)$	40.8%	13.3%	14.2%
	$\sigma\left(R^{stock}_{t,t+1}\right)$	33.5%	42.8%	54.1%
return predictability	$\rho\left(\frac{P_t}{D_t}, R^{stock}_{t,t+4}\right)$	-0.22	-0.07	0.57
	$\rho\left(\frac{P_t}{D_t}, R^{stock}_{t,t+20}\right)$	-0.44	-0.18	0.84

Quarterly US data 1962Q1-2012Q4. Output Y_t is GDP. Consumption C_t consists of of services and non-durable private consumption. Investment I_t consists of private non-residential fixed investment and durable consumption. Wages are total compensation divided by total hours worked. Dividends D_t are four-quarter moving averages of S&P 500 dividends. The stock price index P_t is the S&P500. All variables up to here are deflated by the GDP deflator. Unemployment u_t is the civilian unemployment rate. Returns R_t^{stock} are annualised real quarterly stock returns including dividends. The subscript "hp" indicates application of the Hodrick-Prescott filter with smoothing parameter 1600. Model values are calculated using simulated model data aggregated up to quarterly frequency. † indicates moments targeted by SMM.

Figure 1.5: Impulse responses to a productivity shock in the search model.

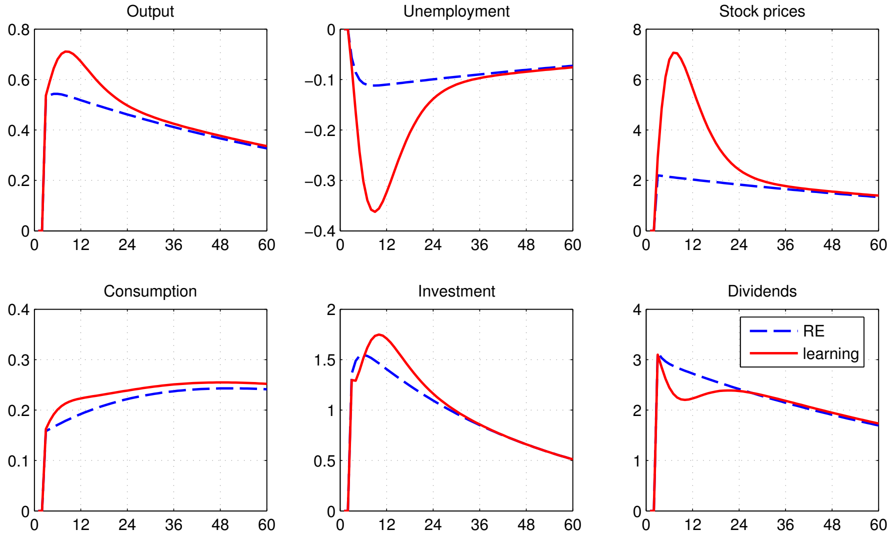

Impulse responses to a one-standard deviation innovation in ε_t. Output Y_t, stock prices P_t, consumption C_t, investment I_t, and dividends D_t are in 100*log deviations from steady state. The unemployment rate u_t is in percentage point deviations from steady state. All responses are at monthly frequency. The RE responses are evaluated at the parameter values for σ_a, η estimated under learning.

solid lines trace responses for the model under learning, while the blue dashed lines trace responses under rational expectations, but at the same parameter values as estimated under learning (i.e. with a low degree of wage rigidity that does not match the volatility of unemployment). This way, the incremental effect of learning should become clearer. The impact response in the first period after the shock is identical under learning and under rational expectations, but after that, stock prices increase markedly under learning due to investors adjusting their beliefs upwards. This leads to an increase in firm value and raises the incentives for entrepreneurs to start new firms and create jobs. The unemployment rate therefore falls by more than under rational expectations. This illustrates why the learning model can replicate the unemployment volatility in the data without relying on a high degree of wage rigidity. By consequence, the response of output is also amplified. Consumption and investment both rise persistently and their response is amplified as well compared to rational expectations.

The response of dividends, however, is smaller under learning than under rational expectations. This can be understood as follows. The expression for dividends is simply

$$d_t = \omega (1-\eta)(1-\alpha)(a_t k_t^\alpha - \bar{y}) + \eta (1-\alpha) \bar{y} \qquad (1.81)$$

where $\bar{y} = \mathbb{E}^{\mathcal{P}}[y_t]$ is the unconditional expectation of y_t. This expression is increasing in the level of capital per firm. At the same time, the optimality condition for the choice of capital can be written as:

$$1 \quad = \quad \mathbb{E}_t^{\mathcal{P}}\left[\beta\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}\left(a_{t+1}k_{t+1}^{\alpha-1}+1-\delta\right)\right] \tag{1.82}$$

Up to first order, the level of capital per firm k_t and dividends d_t are therefore inversely related to expected consumption growth. Under learning, a shock which raises stock prices will raise job creation and therefore increase expected consumption growth. The capital stock per firm will fall as a result, even though the total capital stock of all firms may rise.

This means that the behaviour of dividends dampens rather than reinforces the learning dynamics. The equilibrium stock price is

$$p_t = \frac{d_t}{1-(1-s)\,\mathbb{E}_t^{\mathcal{P}}\left[\beta\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}\right]\exp\left(\hat{\mu}_t+\frac{g-g^2}{2}\sigma_z^2\right)} \tag{1.83}$$

which is increasing in the price growth belief $\hat{\mu}_t$ and increasing in the per-firm dividend d_t. When beliefs $\hat{\mu}_t$ rise, dividends d_t fall over time as a result of decreased capital per firm. Therefore, the amount of amplification in the stock price is lower than if dividends were exogenous.

1.5.4.3 *Wage rigidity and the "Shimer puzzle"*

Table also 1.4 documents that learning reduces the need to rely on wage rigidity to generate the right magnitude of unemployment fluctuations. The estimated value for the wage rigidity parameter is $\omega = 0.134$ under learning, but $\omega = 0.592$ under rational expectations. This means that the amount of wage rigidity required to match the volatility of unemployment is less than four times smaller when learning about stock prices is introduced into the learning. Thus, learning can successfully address the Shimer puzzle successfully while at the same time improving the asset price properties of the search model.

Certainly, the surplus share η of entrepreneurs is still assumed to be very small, and this helps the model generate unemployment volatility as pointed out by Hagedorn and Manovskii (2008). But it is not possible to match unemployment volatility with a low surplus share alone. Figure 1.6 illustrates this. I run the moment matching procedure described above for several values of η both under rational expectations and learning and plot the resulting values of the wage rigidity parameters ω. It is evident that the required value for ω is always substantially lower under learning than under rational expectations. In fact, for values for η above 4%, it is impossible to match unemployment volatility under rational expectations, whereas this is still possible under learning for a value of η up to about 14%.

Figure 1.6: Required degree of wage stickiness.

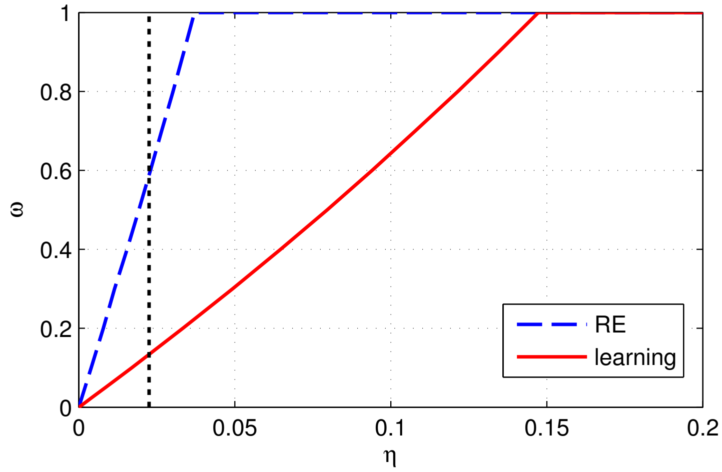

Estimated wage stickiness ω required to match volatility of unemployment, as a function of the surplus share of entrepreneurs η. Values for which moments could not be matched are plotted as $\omega = 1$. Calibration value $\eta = 2.255\%$ marked by black dotted vertical line.

1.6 conclusion

In this chapter, I have discussed the consequences of learning about firm value in two widely used models of the business cycle: the real business cycle model and the labour search and matching model. In doing so, I developed a particular concept of expectation formation called conditionally model-consistent expectations (CMCE), which allow for a subjective belief about one relative price (the stock) price while pinning down the beliefs about all other model variables by requiring a maximum of consistency with model outcomes. This keeps a model as close as possible to rational expectations while still allowing to study the effects of extrapolative expectations about stock prices as in Adam et al. (2013).

It was found that the real business cycle model produces counterfactual results when learning is added. Because a stock price boom mainly causes a large wealth effect, consumption, employment and output tend to move in opposite directions.

By contrast, adding learning to the search and matching model improved the asset price properties of the model while generating positive comovement in all macroeconomic aggregates in response to changes in stock price expectations. What's more, learning substantially reduced the need to rely on wage rigidity to match the magnitude of unemployment fluctuations seen in the data.

The main shortcoming of the analysis in this chapter is that it was not possible to identify a two-sided feedback channel between prices under learning and economic fundamentals. In the RBC model, a stock price boom entailed a surge in dividends but a fall in aggregate output, while in the search model output rose but dividends fell. Are there plausible mechanisms for which the behaviour of dividends and economic activity reinforce belief dynamics under learning? I revisit this question in the next chapter.

The concept of expectations developed in this chapter has many potential applications beyond stock price learning. It allows to study the impact of deviations from rational expectations in one variable while keeping other expectations model-consistent, in virtually any kind of forward-looking model. It would be possible, for example, to apply it to inflation expectations in New Keynesian models, expectations about house prices in models of the housing market, and so on. These applications open up potential avenues for future research.

2 THE ROLE OF LEARNING FOR ASSET PRICES, BUSINESS CYCLES, AND MONETARY POLICY

2.1 INTRODUCTION

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

The above statement by the Chief Economist of the Bank of England at a parliamentary hearing may provoke disbelief among macroeconomists. After all, a wealth of research in the last fifteen years has been dedicated precisely to the links between the financial sector and the real economy. Financial frictions are now seen as a central mechanism by which asset prices interact with macroeconomic dynamics.

Still, our understanding of this interaction remains incomplete, in part due to the inherent difficulty of modelling asset prices.

The typical business cycle model employs an asset pricing theory based on time-separable preferences with moderate degrees of risk aversion and rational expectations. Such an asset pricing theory is well known to be inadequate for many empirical regularities such as return volatility (Shiller, 1981) and return predictability (Fama and French, 1988). This is not problematic when asset prices are disconnected from the real economy, since asset pricing and business cycle dynamics can then be separated.[1] In the presence of financial frictions however, the prices of assets used as collateral affect borrowing constraints and hence the dynamics of the economy. A failure to generate realistic endogenous asset price dynamics can then become a potentially important source of model misspecification.

This chapter examines the business cycle implications of a learning-based asset pricing theory in the presence of financial constraints. I construct a model of firm credit frictions in which agents are unable to form rational expectations about the price of equities in the stock market, and instead have to learn from past observation to form subjective beliefs. The learning-based approach to stock pricing has

1 For the real business cycle model with recursive preferences, Tallarini Jr. (2000) showed that business cycle properties are driven almost entirely by the intertemporal elasticity of substitution, while asset price properties are almost entirely governed by the degree of risk aversion.

been shown to perform surprisingly well in endowment economies, without the need to rely on non-separable preferences or habit (Adam, Marcet and Nicolini, 2013; Adam, Beutel and Marcet, 2014). The interpretation of price dynamics under learning is quite different from rational expectations. With learning, stock prices fluctuate not because of variations in the discounting of prices and returns, but because of variations in subjective beliefs about the prices and returns themselves. The deviation of these subjective beliefs from rational expectations is a natural measure of "price misalignments", "over-" and "undervaluation". These notions are often present in informal arguments about financial markets, but absent in most asset pricing theories.[2]

A second model ingredient is that firms are subject to credit constraints, the tightness of which depends on firm market value. This type of constraint emerges from a limited commitment problem in which defaulting firms can be restructured and resold (similar to Chapter 11 of the US Bankruptcy Code) as opposed to being liquidated. It provides a mechanism by which high stock market valuations translate into easier access to credit. The model has a "financial accelerator" mechanism similar to Bernanke et al. (1999), with the strength and properties of this mechanism crucially depending on the endogenous dynamics of stock prices.

The analysis of the model yields three results. First, a positive feedback loop emerges between beliefs, asset prices and the production side of the economy, which leads to considerable amplification and propagation of business cycle shocks. When investor beliefs are more optimistic, their demand for stocks increases. This increases firm valuations and relaxes credit conditions. This in turn allows firms to move closer to their profit optimum. Provided counteracting general equilibrium forces are not too strong, they will also be able to pay higher dividends to their shareholders, raising stock prices further and propagating investor optimism even more. The financial accelerator mechanism becomes much more powerful than under rational expectations. At the same time, the learning mechanism greatly improves asset price properties such as price and return volatility and predictability without the need to impose complex preferences or high degrees of risk aversion. This result suggests that the relatively weak quantitative strength of the financial accelerator effect in many existing models (Cordoba and Ripoll, 2004) is at least in part due to low endogenous asset price volatility.

Second, while agents' subjective expectations are not rational expectations, they are consistent in a number of ways with data obtained from surveys. I document that forecast errors on several macroeconomic aggregates (from the US Survey of Professional Fore-

2 For example, in October 2014, the IMF warned of "highly correlated mispricing [..] across assets" in its Global Financial Stability Report (p. 6).

casters) as well as on stock returns (from the Duke-Fuqua CFO Survey) can be predicted by the price/dividend ratio as well as forecast revisions. This is also true in the model, despite the fact that learning only adds one free parameter to the model and agents have conditionally model-consistent expectations for all relevant prices and outcomes. When they are over-predicting asset prices, they also over-predict credit limits depending on those prices and therefore aggregate activity, just like in the data.

Third, I show that the model has important normative implications. A recurring question in monetary economics is whether policy should react to asset price "misalignments". Gali (2014) writes that justifying such a reaction requires "the presumption that an increase in interest rates will reduce the size of an asset price bubble" for which "no empirical or theoretical support seems to have been provided". This chapter is a first step towards filling this gap. Indeed, I find that under learning, the welfare-maximising monetary policy within a class of interest rate rules reacts strongly to asset price growth. By raising interest rates when stock prices are rising, policy is able to curb the endogenous build-up of over-optimistic investor beliefs. Such a reaction reduces both asset price volatility and business cycle volatility. In contrast, under rational expectations, a policy reaction to asset prices does not improve welfare, in line with earlier findings in the literature.

The remainder of this chapter is structured as follows. Section 2.2 reviews the related literature. Section 2.3 provides several empirical facts relating to the macroeconomic effects and properties of stock prices as well as discrepancies of measured expectations from rational expectations. Section 2.4 presents a highly stylised version of the model that permits an analytic solution. It shows that credit frictions or asset price learning alone does not generate either amplification of shocks or interesting asset price dynamics, while their combination does. The full model which can be used for quantitative analysis is then presented in Section 2.5. Section 2.6 contains the quantitative results. Section 2.7 contains the monetary policy analysis. Section 2.8 concludes.

2.2 RELATED LITERATURE

This chapter starts from learning-based asset pricing developed in a series of papers by Klaus Adam and Albert Marcet (Adam and Marcet, 2011; Adam et al., 2013, 2014). They show that parsimonious models of learning about stock prices succeed in explaining key aspects of observed stock price data such as the excess volatility, equity premium, and return predictability puzzles. They also show consistency with investor expectations, which are hard to reconcile with rational expectations. Recent work by Barberis et al. (2015) goes in a similar direction. While these papers study endowment economies, I take their approach to an economy with production. This allows to look at the interactions between financial markets and the real economy, as well as policy implications.

There exist other approaches to asset pricing in production economies. For models with financial frictions in particular, it is popular to simply include exogenous shocks to explain the observed fluctuations in asset prices. Iacoviello (2005) and Liu et al. (2013), for example, set up economies in which exogenous shocks to housing demand drive house prices which in turn affect credit constraints, and study the financial accelerator mechanism. Xu et al. (2013) have a model with a credit friction in which borrowing limits also depend on stock market valuations similar to that in my model. They prove the existence of rational liquidity bubbles and introduce a shock that governs the size of this bubble, thus enabling them to match the stock prices seen in the data. In all of these models, the simple preferences and rational expectations would not allow realistic asset price dynamics in the absence of asset price shocks, and the structural interpretation of these shocks is often not clear.[3] In order to advance our understanding of the interaction between asset prices and the real economy, I believe that it is necessary to have macroeconomic models that can endogenise asset price fluctuations. This chapter is a step in this direction.

The macro-finance literature has two main, rational expectations-based propositions to obtain realistic asset price dynamics. The first one is due to Campbell and Cochrane (1999) and relies on a non-linear form of habit formation combined with high risk aversion. The second, so-called "long-run risk" approach due to Bansal and Yaron (2004), introduces small, predictable and observable components to long-run consumption and dividend growth combined with Epstein-Zin preferences. There are some papers that try to embed these alternative approaches in production economies: Boldrin et al. (2001) for habit formation, Tallarini Jr. (2000) and Croce (2014) for long-run risk. These papers consider real business cycle models and are mainly concerned with endogenising consumption and dividend streams in a production economy while preserving the asset price implications. To my knowledge, there are no studies which take either approach to larger business cycle models with financial frictions, possibly because they are computationally quite demanding. They also require the use of preferences which have some rather counter-intuitive properties (shown by Lettau and Uhlig (2000) for habit and Epstein et al. (2013) for long-run risk). In my view, learning-based asset pricing is a promising alternative. It is intuitively appealing to think that asset prices

3 Even the disaster-risk model of Gourio (2012) can be interpreted as such a model in which exogenous shocks to discount factors drive asset prices, even though quantities are not affected by financial frictions but by the changes in discount factors themselves.

are to some degree driven by self-amplifying waves of over- and under-confidence, and such a view is supported by survey evidence. Importantly, it also has implications for policy, as this chapter shows.

The chapter also makes a contribution to the literature on adaptive learning in business cycles. A number of papers in this area have studied learning in combination with financial frictions (Caputo, Medina and Soto, 2010; Milani, 2011; Gelain, Lansing and Mendicino, 2013). The conventional approach taken in this literature consists of two steps: first, derive the linearised equilibrium conditions of the economy under rational expectations; second, replace all terms involving expectations with parametrised forecast functions, and update the parameters using recursive least squares every period. Such models certainly produce very rich dynamics, but they are problematic on several grounds. First, it is not clear that first-order conditions parametrised in this way correspond in a meaningful sense to intertemporal optimisation problems.[4] Second, these models are often very complex and intransparent. The need to parametrise every expectation in the first-order conditions requires a large number of parameters. In all but the simplest models, it then becomes prohibitively difficult to analyse equilibrium dynamics. In this chapter, I make use of a more transparent and parsimonious approach. Beliefs are restricted to be model-consistent as under rational expectations (with the only exception being the beliefs about stock prices) and agents make optimal choices given this set of beliefs. Even in a medium-sized DSGE model, the introduction of learning then adds only one parameter and one state variable to the model.

Finally, the chapter also relates to the debate on whether monetary policy should react to asset price "misalignments". Bernanke and Gertler (2001) found in a financial frictions model with rational exogenous asset price bubbles that the answer is "no". This view, although not unchallenged (Filardo, 2001; Cecchetti et al., 2002), forms the consensus opinion and indeed the practice of most central banks. It has also recently been reinforced by Gali (2014), who argues that since rational bubbles are predicted to grow at the rate of interest, the optimal policy to deflate a bubble might even be to *lower* interest rates when asset prices are rising too fast. Without incorporating bubbles, Faia and Monacelli (2007) find a similar result, and conclude that a strong exclusive anti-inflationary stance remains welfare-maximising. This chapter shows that such policy recommendations depend critically on the underlying asset price theory. In a world of less than fully rational expectations, raising interest rates in an asset price boom can be effective in curbing exuberant investor expectations and mitigate

4 In any model with an Euler equation, for example, some version of the stochastic discount factor needs to be learned by the agents, an object which depends on their own choices. It is unclear how an agent would be able to select a choice based on an intertemporal first-order condition and at the same time not understand how he makes choices when forming expectations.

a surge (and subsequent reversal) in real activity due to high asset prices and easy access to credit.

2.3 EMPIRICAL EVIDENCE

The purpose of this section is to document three sets of facts. First, movements in the aggregate stock market have sizeable effects on investment and credit constraints, consistent with the credit friction in my model. Second, stock prices exhibit high volatility and return predictability. Third, measures of expectations from survey data, both for stock prices and macro variables, reveal systematic deviations from rational expectations. Some but not all of these observations have been documented previously in the literature.

2.3.1 *Effect of the stock market on investment and credit constraints*

One of the oldest documented links between financial markets and the real economy is that the stock market predicts investment (Barro, 1990). Of course, prediction does not imply causation. It is plausible that new information about improved economic fundamentals causes both stock prices and investment to rise, with stock prices responding faster. This view is taken by Beaudry and Portier (2006) who show that in an estimated vector error correction model, innovations in stock prices orthogonal to current changes in TFP predict a substantial portion of long-run TFP variation. This suggests that stock price fluctuations are in fact "news shocks" about future productivity. However, it is also conceivable that stock market movements have a direct effect on investment even when they do not reflect changing expectations about the economic fundamentals. Blanchard et al. (1993) construct a measure of expected fundamentals and find that stock prices retain their predictive power even when controlling for fundamentals. In general though, it is hard to come to any definitive conclusions about causality without spelling out a structural model.

In the model of this chapter, higher stock prices affect investment because of financial frictions: Firms with higher market value have easier access to external finance and can therefore increase investment. Is this consistent with the data? It is, at least when looking at aggregate time series. I estimate a VAR using quarterly US data. The VAR includes six variables: investment, total factor productivity, dividends, the Federal Funds rate, a corporate credit spread, and the aggregate price/dividend-ratio.[5]

5 Investment is real private non-residential fixed investment. Productivity is adjusted for capacity utilisation as in Kimball et al. (2006). Dividends are four-quarter moving averages from the S&P Composite index. The corporate credit spread is Moody's baa-aaa corporate bond spread, serving as a proxy for credit market conditions. The P/D

Figure 2.1: Stock price shock in a VAR.

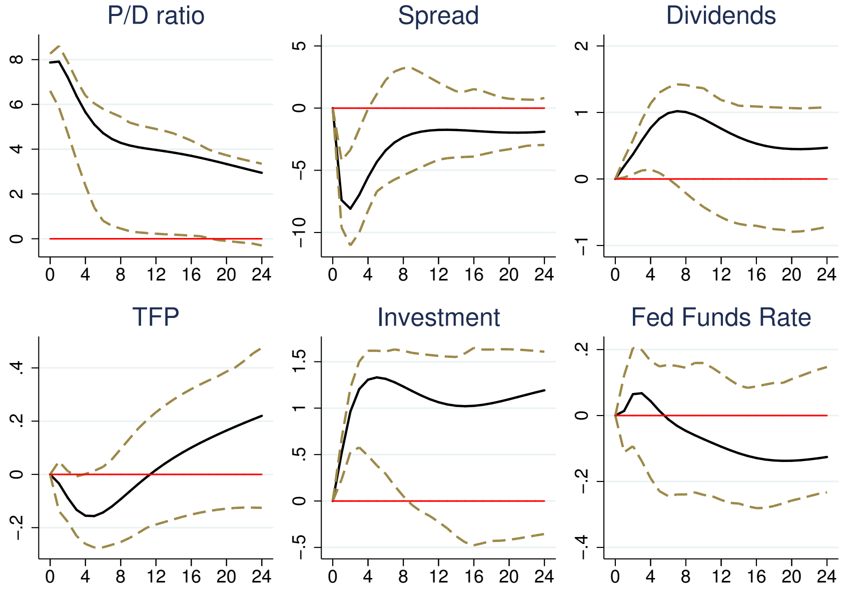

Impulse response functions to a shock to the P/D ratio in a six-variable VAR with lag length 2 (selected by Bayesian Information Criterion) and shock identification by Cholesky decomposition in the displayed ordering. Quarterly US data 1962Q1-2012Q4. Displayed units are basis points for the spread and percentage points for all other series. Bootstrapped confidence bands at the 95% level.

In order to isolate movements in stock prices that are unrelated to contemporaneous productivity, monetary, or other shocks, I examine the effects of a "stock price shock", identified as having an immediate effect on the P/D ratio but no contemporaneous effect on any other variables.[6] This shock alone accounts for more than two thirds of the forecast error variance of the P/D ratio at all horizons.

Figure 2.1 plots the estimated impulse response functions. The shock leads to a persistent rise in the P/D ratio. It also significantly increases investment and dividends while reducing credit spreads. The effect on TFP is insignificant throughout and initially negative. This casts doubt on the view that most stock price movements are a reflection of news about future productivity. They also do not seem to reflect news about interest rates, since the response of the Federal Funds rate, too, is flat and insignificant.

ratio is again from the S&P composite index. The lag length is set to two as per the Bayesian Information Criterion.

6 This ordering is chosen on the premise that financial markets adjust faster to shocks than either real variables or monetary policy. As for the ordering among the financial variables, the results are robust to inverting the order of the P/D ratio and the credit spread.

Table 2.1: Stock market statistics.

	statistic	value
excess volatility	$\sigma\left(\text{hp}\left(\log P_t\right)\right) / \sigma\left(\text{hp}\left(\log D_t\right)\right)$	2.63 (.350)
	$\sigma\left(\log \frac{P_t}{D_t}\right)$	.408 (.017)
	$\sigma\left(\log R^{stock}_{t,t+1}\right)$	.335 (.022)
return predictability	$\rho\left(\frac{P_t}{D_t}, R^{stock}_{t,t+4}\right)$	-.216 (.065)
	$\rho\left(\frac{P_t}{D_t}, R^{stock}_{t,t+20}\right)$	-.439 (.049)

Quarterly data 1949Q1-2012Q4. Stock price and aggregate dividend data for the S&P Composite Index from Robert Shiller's website. Standard errors in parentheses.

2.3.2 *Asset price "puzzles"*

Asset prices in general, and stock prices in particular, are known to exhibit a number of characteristics that are difficult to reconcile with a basic consumption-based asset pricing model (by which I mean a representative investor with time-separable power utility and rational expectations). Here, I document two of them: excess volatility and return predictability.[7] These are summarised in Table 2.1 for quarterly aggregate US data.

The first row shows the ratio of the standard deviation of the cyclical components of stock prices and dividends. By this measure prices are 2.63 times more volatile than dividends. The log price/dividend ratio (second row) and log stock returns (third row) are also highly volatile. Shiller (1981) showed that this amount of volatility cannot be accommodated in an asset pricing theory based on rational expectations and constant discount rates. If one starts from the premise that asset prices equal discounted cash flows, then this implies that either discount rates must vary a lot, or expectations are not rational (or both).

The fourth and fifth rows of the table document return predictability at the one- and five-year horizon, respectively. A high P/D ratio reliably predicts low future returns at these horizons, even if short-run stock returns are almost unpredictable. Cochrane (1992)

7 Another equally famous fact due to Mehra and Prescott (1985) is the size of the equity premium. Adam et al. (2013) show that learning models are able to generate sizeable equity premia, but in this chapter, I only focus on volatility and return predictability.

shows that the variance of the P/D ratio can be decomposed into its covariance with future returns and future dividend growth. Since dividend growth is not very volatile, not well predicted by the P/D ratio, and the P/D ratio itself is volatile, it follows that returns must be predictable. Assuming rational expectations, Cochrane identifies the predictable component of returns with a time-varying discount rate. Again, the alternative is that expectations used to price assets are distinct from rational expectations. In the model of this chapter, a high P/D ratio is not a result of low required returns, but of high expected returns - where the subjective expectation is distinct from the statistical prediction.

2.3.3 *Survey data on expectations*

The rational expectations hypothesis is a fundamental building block of modern macroeconomics. Sometimes, it is criticised as an unrealistic modelling device which asserts that agents are hyper-rational, endowed with infinite computing power and knowledge of the structural shocks and relationships of the economy. But in fact, it makes no such claim. In the words of Sargent (2008), it simply asserts that "outcomes do not differ systematically [...] from what people expect them to be". Put differently, any agent's forecast error should not be predictable by information available to the agent at the time of the forecast.

The rational expectations hypothesis is testable based on survey measures of expectations. It is almost always rejected. Here, I document some of these tests, and characterise some of the predictability patterns of forecast errors.

Expectations of returns are positively correlated with past returns and the P/D-ratio, whereas the best statistical prediction would call for a negative correlation. The difference is strong enough to be statistically rejected. This pattern is observable across many different sources of survey data (Greenwood and Shleifer, 2014). I illustrate it with data from the CFO survey by John Graham and Campbell Harvey at Duke University. The survey respondents are CFOs of major US corporations, which are likely to possess good knowledge of financial markets. Since 2000, the survey includes a question on stock market return expectations ("Over the next year, I expect the average annual S&P 500 return will be ..."). Figure 2.2 compares the survey expectations with realised returns. The left panel plots the mean survey response against the value of the P/D ratio in the month preceding the survey. The correlation is strongly positive: return expectations are more optimistic when stock valuations are high. However, high stock valuations actually predict low future returns, as documented above and illustrated again in the right panel of the figure. Such a pattern cannot be reconciled with rational expectations, as it implies

Figure 2.2: Return expectations and expected returns.

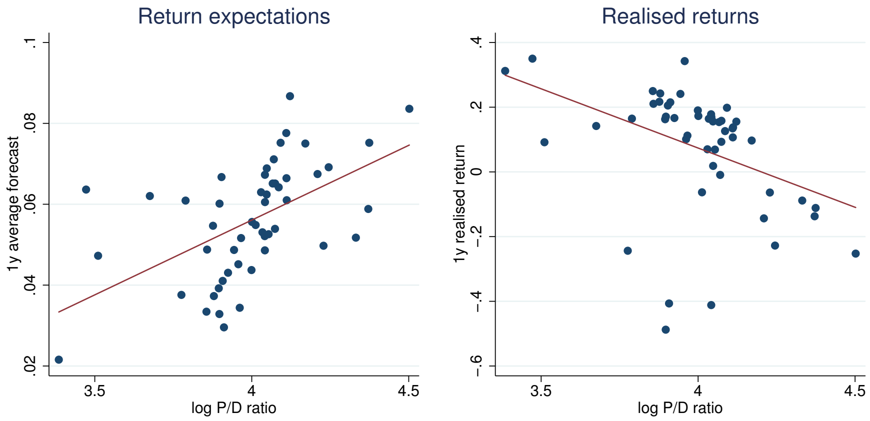

Expected nominal returns are the mean response in the Graham-Harvey survey, P/D ratio and realised nominal returns data from the S&P 500. Data period 2000Q3-2012Q4. Correlation coefficient for expected returns $\rho = .54$, for realised returns $\rho = -.40$.

that agents' forecast errors are predictable by the P/D ratio, a publicly observable statistic.[8]

Of course, survey data are only an imperfect measure of expectations. When answering questions on a survey, respondents might wilfully misstate their true expectations, answer carelessly, or misunderstand the question. When being asked for a point estimate, they might report a statistic other than the mean of their belief distribution. Still, survey data are the best available test for the rational expectations hypothesis.

Tests of forecast error predictability can be applied to other variables of macroeconomic significance. Table 2.2 describes tests using the Federal Reserve's Survey of Professional Forecasters (SPF) as well as the CFO survey data. Each row and column corresponds to a correlation of a mean forecast error with a variable that is observable by respondents at the time the survey is conducted. Under the null of rational expectations, the true correlation coefficients should all be zero.

The first column shows that the P/D ratio negatively predicts with forecast errors. When stock prices are high, people systematically under-predict economic outcomes. This holds in particular for stock returns, as was already shown in the scatter plot above. But it also holds true for macroeconomic aggregates, albeit at lower levels of significance. The second column shows that the change in the P/D ratio

8 To be precise, one can test the hypothesis that the univariate regression coefficient of the P/D ratio on expected (survey) returns is the same as that on realised returns, e.g. using the SUR estimator. The hypothesis is rejected at the 0.1% level. The survey asks for a return estimate of a 12-month period starting at varyind days of the month, whereas the realised return measure is taken at the beginning of the month. But the results are robust to taking realised returns at the end of the month.

Table 2.2: Forecast error predictability: Correlation coefficients.

forecast variable	(1) log PD_t	(2) $\Delta \log PD_t$	(3) forecast revision
$R^{stock}_{t,t+4}$	-.44***	.06	n/a
	(-3.42)	(.41)	
Y_{t+3}	-.21*	.22**	.29***
	(-1.78)	(2.42)	(3.83)
I_{t+3}	-.20*	.25***	.31***
	(-1.74)	(2.88)	(3.79)
C_{t+3}	-.19*	.21**	.23***
	(-1.85)	(2.37)	(2.67)
u_{t+3}	.05	-.27***	.43***
	(.12)	(-3.07)	(6.07)

Correlation coefficients for mean forecast errors on one-year ahead nominal stock returns (Graham-Harvey survey) and three-quarter ahead real output growth, investment growth, consumption growth and the unemployment rate (SPF). t-statistics for the null of zero correlation in parentheses. One, two, and three asterisks correspond to significance at the 10%, 5%, and 1% levels. 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 (2010). Data from Graham-Harvey covers 2000Q3-2012Q4. Data for the SPF covers 1981Q1-2012Q4.

is a better predictor of forecast errors, strongly rejecting the rational expectations hypothesis, but in the opposite direction. At times when stock prices are rising, people systematically under-predict economic outcomes. Since the stock market itself positively predicts economic activity (as shown above), this suggests that agents' expectations are too cautious and under-predict an expansion in its beginning, but then overshoot and over-predict it when it is about to end. Such a pattern emerges naturally in the model of this chapter under learning.

The third column reports the results of a particular test of rational expectations devised by Coibion and Gorodnichenko (2010). Since for any variable x_t, the SPF asks for forecasts at one- through four-quarter horizon, it is possible to construct a measure of agent's revision of the change in x_t as $\hat{\mathbb{E}}_t \left[x_{t+3} - x_t\right] - \hat{\mathbb{E}}_{t-1} \left[x_{t+3} - x_t\right]$. Forecast errors are positively predicted by this revision measure. Coibion and Gorodnichenko take this as evidence for sticky information models in which information sets are gradually updated over time. As I will show later, it is also consistent with the learning model developed in this chapter.

2.4 UNDERSTANDING THE MECHANISM

In this section, I construct a simplified version of the model which illustrates the interaction between asset prices and credit frictions under learning. I impose several strong assumptions permitting a closed-form solution. Quantitative analysis will require a richer model, the development of which is relegated to the next section. The main message of this section is that financial frictions alone do not generate either sizeable amplification of business cycle shocks or asset price volatility, but in combination with learning they do.

2.4.1 *Model setup*

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

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

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

C_t is the amount of non-durable consumption goods purchased by the household in period t. The consumption good is traded in a competitive spot market and serves as the numÈraire. w_t is the real wage rate. Labour is also traded in a competitive spot market, so that the wage w_t is taken as given by the household. Moreover, the household can trade two financial assets, again in competitive spot markets: one-period bonds, denoted by B_t and paying gross real interest R_t in the next period; and stocks S_t which trade at price P_t and entitle their holder to dividend payments D_t. The household cannot short-sell stocks and his maximum stock holdings are capped at some $\bar{S} > 1$.[9]

The household maximises the expectation of discounted future consumption under the probability measure $\mathcal{P}$. This measure is the subjective belief system held by agents in the model economy at time t, which will be discussed in detail further below. The first order conditions describing the household's optimal plan under an arbitrary $\mathcal{P}$ are

$$R_t = R = \beta^{-1} \tag{2.1}$$

9 The constraint on S_t is necessary to guarantee existence of the learning equilibrium, although it never binds along the equilibrium path.

$$S_t \begin{cases} = 0 & \text{if } P_t > \beta \mathbb{E}_t^{\mathcal{P}} \left[P_{t+1} + D_{t+1} \right] \\ \in [0, \bar{S}] & \text{if } P_t = \beta \mathbb{E}_t^{\mathcal{P}} \left[P_{t+1} + D_{t+1} \right] \cdot \\ = \bar{S} & \text{if } P_t < \beta \mathbb{E}_t^{\mathcal{P}} \left[P_{t+1} + D_{t+1} \right] \end{cases} \tag{2.2}$$

Let us now turn to the firm. It engages in the production of a good which can be used both for consumption and investment. It is produced using capital K_{t-1}, which the firm owns, and labour L_t according to the constant returns to scale technology

$$Y_t = K_{t-1}^{\alpha} \left(A_t L_t \right)^{1-\alpha} \tag{2.3}$$

where A_t is its productivity. Here, I only allow for permanent shocks to productivity:

$$\log A_t = \log G + \log A_{t-1} + \varepsilon_t,\ \varepsilon_t \sim \mathcal{N} \left(-\frac{\sigma^2}{2}, \sigma^2 \right) \ iid \tag{2.4}$$

In particular, the expected growth rate of productivity is constant at $\mathbb{E}_t A_{t+1} / A_t \equiv G$. The capital stock is predetermined and owned by the firm. It depreciates at the rate δ at the end of each period. The firm can also issue shares and bonds as described above. Thus, its period budget constraint reads as follows:

$$Y_t + (1 - \delta) K_{t-1} + B_t + S_t P_t = w_t L_t + K_t + S_{t-1} \left(P_t + D_t \right) + R B_{t-1} \tag{2.5}$$

Before describing the equilibrium, it is useful to introduce the marginal return on capital:

$$R_t^k = \frac{\partial Y_t}{\partial K_{t-1}} + 1 - \delta \tag{2.6}$$

2.4.2 *Frictionless equilibrium*

In the absence of financial constraints, the Modigliani-Miller theorem will render the composition of firm financing redundant and the model collapses to a standard stochastic growth model. In particular, the optimal choice of the capital stock (under rational expectations) equates the marginal return on capital with the inverse of the discount factor: $\mathbb{E}_t R_{t+1}^k = \beta^{-1}$. Whatever one assumes about the financial structure of the firm, its total value will equal the size of the capital stock. The capital stock and firm value co-moves perfectly with productivity:

$$K_t / A_t = \tilde{K}_t \quad = \quad K^* \tag{2.7}$$

$$\frac{P_t + B_t}{A_t} = \tilde{K}_t \quad = \quad K^* \tag{2.8}$$

where $K^* = G \left(\frac{\alpha}{\beta^{-1} - 1 + \delta} \right)^{1/(1-\alpha)} e^{-\alpha \sigma^2 / 2}$.

2.4.3 *Rational expectations equilibrium with financial frictions*

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

$$D_t = Y_t - w_t L_t - \delta K_{t-1} - (R-1)\, B_{t-1} \tag{2.9}$$

Combined with equation (2.5), this assumption implies that the firm's capital stock must be entirely debt-financed: $K_t = B_t$ at all t.[10] In other words, the firm's book value of equity after dividend payouts is constrained to be zero. This does not mean, however, that the market value of equity is also zero as long as the firm's expected dividend payouts are strictly positive.

On the debt side, the level of debt that can be acquired by the firm is limited to a fraction $\xi \in [0,1]$ of the total market value of its assets, i.e. the sum of debt and equity:

$$\begin{aligned} B_t &\leq \xi\,(B_t + P_t) \\ \Leftrightarrow K_t &\leq \frac{\xi}{1-\xi} P_t \end{aligned} \tag{2.10}$$

Equation (2.10) is a simple constraint on leverage, i.e. debt divided by value of total assets. I depart from the standard assumption that total assets enter with their liquidation value (in this case K_t, the book value) and instead let them enter with their market value (in this case $B_t + P_t$). This captures the idea that a firm which is more highly valued by financial markets will have easier access to credit. This could be because high market value acts as a signal to lenders for the firm's ability to repay, or because the amount lenders can recover in the event of default depends on the price at which a firm can be resold to other financial market participants. In the full version of the model in the next section, I formally derive (2.10) from a limited commitment problem.

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

$$\max_{(K_t, L_t, D_t)_{t=0}^{\infty}} \mathbb{E} \sum_{t=0}^{\infty} \beta^t D_t$$

$$\begin{aligned} \text{s.t. } D_t &= Y_t - w_t L_t + (1-\delta-R)\, K_{t-1} \\ K_t &\leq \frac{\xi}{1-\xi} P_t \\ K_t &\geq 0,\ K_{-1} \end{aligned}$$

10 This holds under the suitable initial condition $K_{-1} = B_{-1}$, e.g. the firm starts with zero book value of equity on its balance sheet. This assumption is relaxed in the full model.

In particular, it makes its decisions under the same belief system $\mathcal{P}_t$ as the household. Due to constant returns to scale in production, we can write dividends at the optimum as follows:

$$\max_{L_t} D_t = \left(R_t^k - R\right) K_{t-1} \tag{2.11}$$

The optimal choice of capital is to exhaust borrowing limits as long as the expected internal return on capital exceeds the external return paid to creditors:

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

When solving for the equilibrium, market clearing needs to be imposed. The market clearing condition for bonds is just $R = \beta^{-1}$. That for equity is $S_t = 1$, which means the Euler equation (2.2) has to hold with equality:

$$P_t = \beta \mathbb{E}_t \left[P_{t+1} + D_{t+1}\right] \tag{2.13}$$

Goods market clearing requires

$$Y_t + K_t = C + (1-\delta) K_{t-1} \tag{2.14}$$

and labour market clearing requires $L_t = 1$.

The equilibrium under rational expectations admits a closed-form solution. First, note that the equilibrium return on capital depends on the aggregate capital stock due to decreasing returns to scale at the aggregate level:

$$R_t^k = R^k\left(\tilde{K}_{t-1}, \varepsilon_t\right) = \alpha \left(\frac{Ge^{\varepsilon_t}}{\tilde{K}_{t-1}}\right)^{1-\alpha} + 1 - \delta \tag{2.15}$$

This dependency comes about through a general equilibrium effect: A higher level of the capital stock increases labour demand, which increases real wages and therefore lowers the return on capital. Next, one can write expected dividends as a function of the capital stock:

$$\tilde{D}\left(\tilde{k}_t, \tilde{K}_t\right) = \mathbb{E}_t\left[\frac{D_{t+1}}{A_t}\right] = \left(\mathbb{E}_t R^k\left(\tilde{K}_t, \varepsilon_{t+1}\right) - R\right) \tilde{k}_t \tag{2.16}$$

Here, I have made a distinction between the capital choice $\tilde{k}_t$ of the representative firm that takes future wages as given, and the aggregate capital stock $\tilde{K}_t$ which determines wages and the return on capital in general equilibrium. Of course, in equilibrium the two are equal. Finally, the effective stock price under rational expectations is simply the discounted sum of future dividends:

$$\tilde{P}_t = \frac{P_t}{A_t} = \beta \mathbb{E}_t \sum_{s=0}^{\infty} \beta^s \frac{A_{t+s}}{A_t} \tilde{D}\left(\tilde{K}_{t+s}, \tilde{K}_{t+s}\right) \tag{2.17}$$

Figure 2.3: Rational expectations equilibrium.

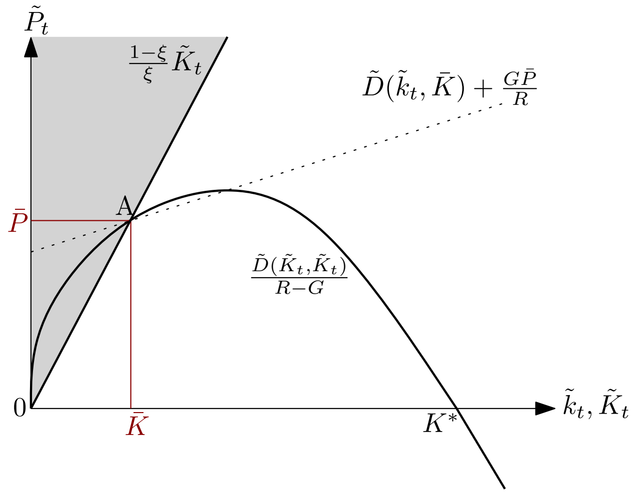

First, let us consider the case $\xi = 1$. In this case, $\tilde{K}_t$ is constant across time and states in equilibrium and is at its efficient level K^* such that $\mathbb{E}_t R^k \left(K^*, \varepsilon_{t+1}\right) = R$. By consequence, expected dividends and the market value of equity are zero: $\tilde{D}\left(K^*, K^*\right) = 0$ and $P_t = 0$. Intuitively, when the firm can borrow up to the total amount of its market value, it faces no financial friction. Book and market value of the firm coincide. The expected return on capital equals the interest rate due to risk neutrality of households, and since all capital is financed by debt and the production function has constant returns to scale, in expectation all profits are paid out as interest payments to debt holders. The residual equity claims trade at a price of zero. This result is the expectation version of the zero-profit condition under constant returns to scale.

Next, let us turn to the case in which ξ is strictly smaller than one. The equilibrium effective capital stock $\tilde{K}_t$ and stock price $\tilde{P}_t$ turn out to be constant here as well. The equilibrium is characterised by two equations:

$$\bar{P} = \frac{\tilde{D}\left(\bar{K}, \bar{K}\right)}{R - G} \tag{2.18}$$

$$\bar{K} = \frac{\xi}{1-\xi}\bar{P} \tag{2.19}$$

The first equation pins down the stock market value of the firm as a function of its capital stock, while the second determines the capital stock that can be reached by exhausting the borrowing constraint that depends on the stock market value. In particular, the internal rate of return is always greater than the return on debt and so the borrowing constraint is always binding. The equilibrium is depicted graphically in Figure 2.3. Equation (2.19) is represented by a straight line of slope $\left(1-\xi\right)/\xi$, while the market value of equity (2.18) is a hump-shaped curve. The equilibrium lies at the intersection (Point A). Firms pay positive dividends (in expectation) and the market value of equity is positive. The capital stock $\bar{K}$ remains inefficiently low: $\bar{K} < K^*$. In addition, Figure 2.3 plots as a dotted line firm value as a function of the current choice of capital $\tilde{k}_t$. From the firm perspective (taking factor prices as given), the value function is increasing everywhere, so it always wants to exhaust the borrowing constraint.

While the capital stock and expected output are an increasing function of maximum leverage ξ, expected dividends are non-monotonous and hump-shaped. Why is that so? There are two opposing forces affecting expected dividends, as can be seen from the following decomposition:

$$\frac{d}{d\xi}\tilde{D}\left(\bar{K},\bar{K}\right) = \left[\underbrace{\mathbb{E}_t R^k\left(\bar{K},\varepsilon_{t+1}\right) - R}_{>0} + \mathbb{E}_t \underbrace{\frac{dR^k}{d\tilde{K}}\left(\bar{K},\varepsilon_{t+1}\right)\bar{K}}_{<0}\right]\underbrace{\frac{d\bar{K}}{d\xi}}_{>0} \tag{2.20}$$

The first term in brackets captures a partial equilibrium effect, which is internalised by the firm. When a firm is financially constrained, its internal rate of return is higher than the return it has to pay to debt holders. By borrowing more, it can increase its scale of production and make more profit. The second term, however, captures a general equilibrium effect: Higher investment lowers the marginal product of capital, which in practice is realised through an increase in the equilibrium wage w_{t+1}. When ξ is small (financial frictions are severe) the partial equilibrium effect dominates, while for a large ξ the general equilibrium effect dominates.

Most importantly however, 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. This can be seen by looking at the variances of log stock price and output growth which do not depend on ξ:

$$\mathrm{Var}\left[\Delta \log P_t\right] = \sigma^2 \tag{2.21}$$

$$\mathrm{Var}\left[\Delta \log Y_t\right] = \left(1 - 2\alpha + 2\alpha^2\right)\sigma^2 \tag{2.22}$$

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, the model cannot replicate many of the stylised facts on stock price data. Up to a first-order approximation, the relative volatility of asset price growth with respect to dividend growth is bounded from below:

$$\frac{\sigma\left(\Delta \log P_t\right)}{\sigma\left(\Delta \log D_t\right)} < \left(1 - 2\alpha + 2\alpha^2\right)^{1/2} \leq \sqrt{2} \tag{2.23}$$

The asset price volatility observed in the data can therefore not be matched. The volatility of the price/dividend ratio can also not be matched: in fact, the forward P/D ratio is even constant:

$$PD_t = \frac{P_t}{\mathbb{E}_t D_{t+1}} \equiv \frac{1}{R-G} \tag{2.24}$$

Furthermore, excess returns are unpredictable: $\mathbb{E}_t\left[\left(P_{t+1}+D_{t+1}\right)/P_t\right] - R = 0$. Finally, by definition of rational expectations, forecast errors are unpredictable, again at odds with the data.

2.4.4 *Learning equilibrium*

I now describe the equilibrium under learning. Qualitatively, this alteration to the expectation formation process can address many of the issues encountered in the last section: it will increase the volatility of stock prices, account for return and forecast error predictability, and most importantly, induce endogenous amplification and propagation on the production side of the model economy. The only departure from rational expectations will be that agents do not understand the pricing function that maps fundamentals into an equilibrium stock price. Instead, they form subjective beliefs about the law of motion of prices and update them using realised price observations.

More specifically, agents continue to make optimal choices for a consistent belief system governed by the measure $\mathcal{P}$. The equilibrium therefore satisfies "internal rationality" as formalised by Adam and Marcet (2011) and used in Eusepi and Preston (2011). I impose the following restrictions to beliefs. Under $\mathcal{P}$,

1. agents observe the exogenous productivity shock ε_t;

2. agents believe that P_t evolves according to

$$\log P_t - \log P_{t-1} = \mu_t + \eta_t \tag{2.25}$$

$$\mu_t = \mu_{t-1} + \nu_t \tag{2.26}$$

$$\text{where } \begin{pmatrix} \eta_t \\ \nu_t \end{pmatrix} \sim \mathcal{N}\left(-\frac{1}{2}\begin{pmatrix} \sigma_\eta^2 \\ \sigma_\nu^2 \end{pmatrix}, \begin{pmatrix} \sigma_\eta^2 & 0 \\ 0 & \sigma_\nu^2 \end{pmatrix}\right) \; iid, \tag{2.27}$$

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

$$\mu_0 \mid \mathcal{F}_0 \sim \mathcal{N}\left(\hat{\mu}_0, \sigma_\mu^2\right) \text{ where } \sigma_\mu^2 = \frac{-\sigma_\nu^2 + \sqrt{\sigma_\nu^4 + 4\sigma_\nu^2\sigma_\eta^2}}{2}; \tag{2.28}$$

3. agents update their beliefs about μ_t after making their choices and observing equilibrium prices in period t;

4. for any variable x_t other than the stock price, any future date $t+\tau$, and any sequence of exogenous productivity shocks $\varepsilon_t, \ldots, \varepsilon_{t+\tau}$ and stock prices $P_t, \ldots, P_{t+\tau}$ which is on the equilibrium path:,

$$\mathbb{E}_t^{\mathcal{P}}\left[x_{t+\tau} \mid \varepsilon_t, P_t, \ldots, \varepsilon_{t+\tau}, P_{t+\tau}\right] = \mathbb{E}_t\left[x_{t+\tau} \mid \varepsilon_t, P_t, \ldots, \varepsilon_{t+\tau}, P_{t+\tau}\right]$$

   i.e. agents' beliefs coincide with the best statistical prediction of x_t *conditional* on the realisation of stock prices and fundamentals.

The first assumption implies that agents have as much information about the fundamental shocks of the economy as under rational expectations. The second assumption amounts to saying that agents believe stock prices to be a random walk. This random walk is believed to have a small, unobservable, and time-varying drift μ_t. Learning about this drift is going to be the key driver of asset price dynamics. The third assumption imposes that forecasts of stock prices are updated after equilibrium prices are determined, so as to avoid possible multiple equilibria in price and forecast determination.[11] The final assumption amounts to imposing conditionally model-consistent expectations (CMCE) developed in Chapter 1.

Despite such relatively accurate beliefs, agents are still left with the important problem of making a good guess about the unobservable drift μ_t of stock prices. Optimal Bayesian belief updating amounts to Kalman filtering in this case, since (2.25)-(2.28) is a linear state-space system. Under $\mathcal{P}$, agents' beliefs about μ_t at time t are normally distributed with stationary variance σ_μ^2 and mean $\hat{\mu}_{t-1}$. This belief about the mean of μ_t, which I will usually just call "the belief", evolves according to the updating equation:

$$\hat{\mu}_t = \hat{\mu}_{t-1} - \frac{\sigma_\nu^2}{2} + g\left(\log P_t - \log P_{t-1} + \frac{\sigma_\eta^2 + \sigma_\nu^2}{2} - \hat{\mu}_{t-1}\right) \tag{2.29}$$

In this equation, P_t and P_{t-1} are observed, realised stock prices. These are determined in equilibrium under the actual law of motion of the economy and do *not* follow the perceived law of motion described by (2.25)-(2.28). The parameter g is called the "learning gain". It governs the speed with which agents move their prior in the direction of the last forecast error.[12] When g is high, agents are confident that observed changes in the growth rate of asset prices are due to changes in the trend μ_t rather than the noise η_t. The gain is not decreasing

11 This "lagged belief updating" is common in the learning literature. It makes all feedback between forecasts and prices inter- rather than intramporal. For further discussion see Adam et al. (2014).

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

in time: Agents believe that the drift in asset prices is itself time-varying, so that even after a long period of time it remains difficult to forecast it. A consequence of this is that beliefs never converge: Agents always entertain the possibility of some structural change to the law of motion of asset prices, so that even an infinite number of observations is not completely informative about the future.

It is important to keep in mind that the disturbances η_t and ν_t are objects that exist only in the subjective belief system $\mathcal{P}$. The actual equilibrium under learning does not contain any shock process other than the productivity shock ε_t. Instead, the equilibrium still contains the actual market clearing conditions (2.13) and (2.14), even if they are unknown to the agents. By Walras' law, it is sufficient to impose stock market clearing. Under $\mathcal{P}$, (2.13) reads as follows:

$$P_t = \frac{\mathbb{E}_t^{\mathcal{P}} P_{t+1} + \mathbb{E}_t^{\mathcal{P}} D_{t+1}}{R} \tag{2.30}$$

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

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

The second line is obtained by substituting in agents' beliefs about the evolution of the future stock price $\mathbb{E}_t P_{t+1}$ and dividends $\mathbb{E}_t D_{t+1}$. Under $\mathcal{P}$, agents forecast future dividends accurately conditional on their belief about stock prices. Therefore, their expectations about dividends depend on the current capital stock in the same way as under rational expectations.

In sum, the learning equilibrium is the solution to the following:

$$\tilde{P}_t = \frac{\tilde{D}\left(\tilde{K}_t, \tilde{K}_t\right)}{R - \exp\left(\hat{\mu}_{t-1} + \frac{1}{2}\sigma_\mu^2\right)} \tag{2.33}$$

$$\tilde{K}_t = \frac{\xi}{1-\xi}\tilde{P}_t \tag{2.34}$$

$$\hat{\mu}_{t+1} = \hat{\mu}_t - \frac{\sigma_\nu^2}{2} + g\left(\log\frac{\tilde{P}_t}{\tilde{P}_{t-1}} + \log G + \varepsilon_t - \hat{\mu}_t + \frac{\sigma_\eta^2 + \sigma_\nu^2}{2}\right) \tag{2.35}$$

The first two equations are static and the third is dynamic. The third equation also depends on the productivity innovation ε_t, and as such $\tilde{P}_t$ and $\tilde{K}_t$ are not constant any more. The resulting stock price dynamics after a positive innovation at $t = 1$ are depicted in Figure 2.4.[13] The initial shock at $t = 1$ raises stock prices proportional to productivity. In the rational expectations equilibrium, this would be all that happens. But learning investors are not sure whether the rise

13 The figure depicts the case in which beliefs start at their rational expectations value $\hat{\mu}_t = \log G$ and subjective uncertainty is vanishing in the sense that $\sigma_\nu^2, \sigma_\eta^2 \to 0$ while g remains constant.

Figure 2.4: Stock price dynamics under learning.

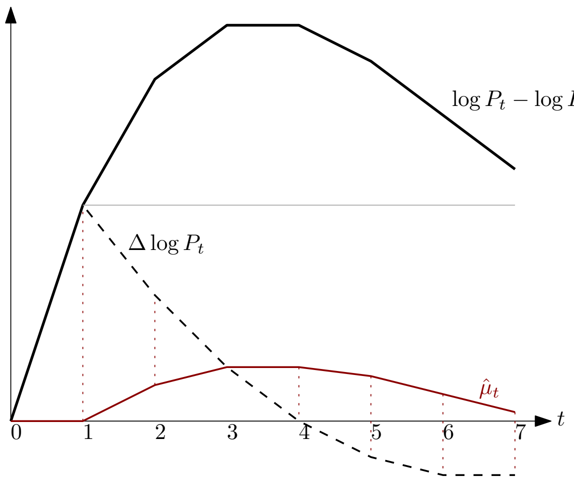

in P_1 is indicative of a transitive or permanent increase in the growth rate of stock prices. They therefore revise their beliefs upwards. In $t = 2$ then, demand for stocks is higher and stock prices need to rise further to clear the market. Beliefs continue to rise as long as observed asset price growth (dashed black line in Figure 2.4) is higher than the current belief $\hat{\mu}_t$ (solid red line). The differences between observed and expected price growth are the forecast errors (dotted red lines). In the figure, the increase in prices and beliefs ends at $t = 3$, when the forecast error is zero. There is no need for a further belief revision. Now, by equations (2.33)-(2.34), in the absence of subsequent shocks, P_t just co-moves with beliefs $\hat{\mu}_{t-1}$, so when there is no belief revision at $t = 4$, realised asset price growth is also zero. This triggers an endogenous reversal in prices, as investors observe stalling asset prices at the peak of their optimism. They subsequently revise their beliefs $\hat{\mu}_t$ downwards, pushing the stock price down until it returns to its steady-state level.

Figure 2.4 can be used to illustrate the key properties of beliefs and prices under learning. First, the described dynamics do not depend on time: The system of $\tilde{K}_t$, $\tilde{P}_t$, and $\hat{\mu}_t$ is stationary. Second, asset prices are more volatile than under rational expectations. Third, (excess) returns are predictable even though agents are risk-neutral and the discount factor is constant. To see this, it is again convenient to look at the forward P/D ratio:

$$PD_t = \frac{1}{R - \exp\left(\hat{\mu}_{t-1} + \frac{1}{2}\sigma_{\mu}^2\right)}$$

The forward P/D ratio is directly related to the belief $\hat{\mu}_t$. A high P/D ratio is realised at the asset price peak in periods 3 and 4 of Figure 2.4, immediately after which stock prices start declining. Therefore, the P/D ratio negatively predicts future returns. Furthermore, forecast errors are also predictable: By Equation (2.35), forecast errors are a

Figure 2.5: Endogenous response of dividends.

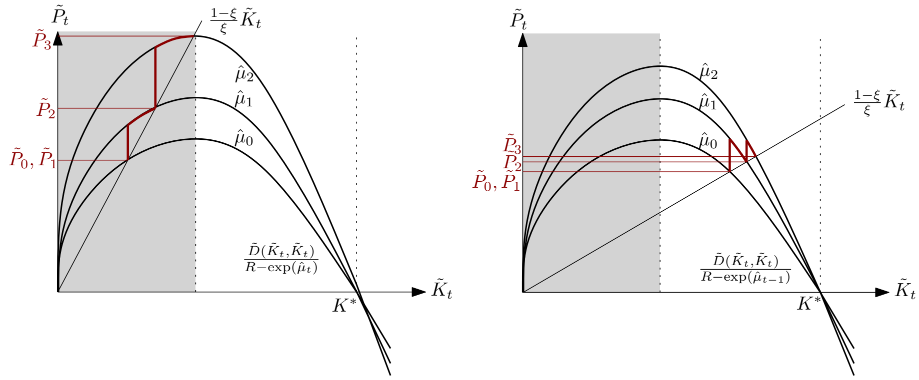

(a) Amplification. (b) Dampening.

linear function of $\hat{\mu}_t - \hat{\mu}_{t-1}$. Since a high P/D ratio implies declining beliefs in the future, forecast errors are predictable in the same way as future returns (as in the data).

The aforementioned asset pricing implications are present even when dividends are completely exogenous, as in Adam et al. (2013). But the model considered here also contains a link between asset prices, output and dividends. The effective capital stock $\tilde{K}_t$ is directly related to equity valuations $\tilde{P}_t$ through Equation (2.34). Thus, the fluctuations in the stock market translate into corresponding fluctuations in investment, the capital stock and hence output. Thus, the presence of financial frictions in combination with learning leads to amplification of the productivity shock, whereby under rational expectations amplification was zero.

It is also possible that this amplification mechanism is *further* enhanced by positive feedback from capital to expected dividends. This additional feedback, however, depends on the slope of the function $\tilde{D}$. As demonstrated above, the expected dividend $\tilde{D}\left(\tilde{k}_t, \tilde{K}_t\right)$ is increasing in the firm's capital choice $\tilde{k}_t$, but decreasing in the aggregate capital stock $\tilde{K}_t$. In the equilibrium ($\tilde{k}_t = \tilde{K}_t$) it is increasing if financial frictions are sufficiently severe. This case is depicted in Panel (a) of Figure 2.5. When the degree of financial frictions is high, the credit constraint line is steep. Assume that the initial equilibrium in period 0 is at $\tilde{P}_0$ and $\hat{\mu}_0$. Now consider the effect of a positive productivity shock in period 1 as before. The immediate effect will be a proportionate rise in stock prices and capital which leaves $\tilde{P}_1$ and $\tilde{K}_1$ unchanged, but raises beliefs from $\hat{\mu}_0$ to $\hat{\mu}_1$. This leads to higher stock prices at $t = 2$ and allows the firm to invest more and increase its expected profits $\tilde{D}\left(\tilde{K}_2, \tilde{K}_2\right)$. But this adds further to the rise in realised stock prices, further relaxing the borrowing constraint and increasing next period's beliefs. Stock prices, beliefs, investment, and output all rise by more compared to a situation in which $\tilde{D}$ is constant.[14]

However, this additional amplification channel only works when ξ is sufficiently low. In Panel (b), ξ is large and the firm is operating in the downward-sloping bit of the profit curve $\tilde{D}$. A relaxation of the borrowing constraint due to a rise in $\hat{\mu}$ still allows the firm to invest and produce more, but dividends fall in equilibrium. This is due to the general equilibrium forces mentioned earlier: The marginal product of capital has to fall in equilibrium, which in practice derives from an increase in real wages, effectively reducing the firm's profits. In this situation, the endogenous response of dividends dampens rather than amplifies the dynamics of investment and asset prices.

This can also be seen algebraically. Under rational expectations, the derivative of log stock prices with respect to the productivity shock is simply $d\log P_t / d\varepsilon_{t-s} = 1$ for all $t, s \geq 0$. With learning, the corresponding expression contains additional terms:

$$\frac{d\log P_t}{d\varepsilon_{t-s}} = 1 + \frac{1}{1 - \epsilon_D\left(\tilde{K}_t\right)} \frac{e^{\hat{\mu}_{t-1}}}{R - e^{\hat{\mu}_{t-1}}} \frac{d\hat{\mu}_t}{d\varepsilon_{t-s}} \tag{2.36}$$

where $\epsilon_D\left(\tilde{K}_t\right) = \frac{d\tilde{D}}{d\tilde{K}}\frac{\tilde{K}}{\tilde{D}}$ is the elasticity of expected dividends with respect to the capital stock. Learning adds a product of three terms. The last term is the effect of the productivity shock on subjective beliefs. The second term is the effect of beliefs on stock prices. Small variations in beliefs $\hat{\mu}_t$ cause large fluctuations in stock prices because the denominator $\left(R - e^{\hat{\mu}_{t-1}}\right)$ is close to zero. The first term captures the general equilibrium effects mentioned earlier: when dividends rise after a relaxation of credit constraints, ϵ_k^d is positive and the term is greater than one, leading to additional amplification, and to dampening when ϵ_k^d is negative.

It can also be shown that the learning dynamics vanish as the economy approaches the unconstrained first-best:

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

In other words, amplification rests on the interaction between learning and financial frictions, not on either of them separately. Intuitively, as financial frictions disappear, the economy moves into a region where the general equilibrium effects become so strong that any potential rise in beliefs or asset prices is countered by a fall in expected dividends.

---

14 To my knowledge, this chapter is the first to establish a positive feedback from fundamentals to beliefs under learning. Adam et al. (2012) also model economies with endogenous fundamentals. Their learning specification is similar, but the "dividend" in their asset pricing equation is simply the marginal utility of housing which is strictly decreasing in the level of the housing stock. Their model dynamics are therefore always as in case (b) described above.

In sum, the learning equilibrium can qualitatively account for a number of asset pricing facts and for the predictability of forecast errors. At the same time, the larger endogenous asset price volatility induces corresponding fluctuations in the slackness of the firm's borrowing constraint. The presence of financial frictions thus magnifies productivity shocks, while there is no amplification under rational expectations. When financial frictions are sufficiently severe, a two-sided positive feedback loop emerges between beliefs, asset prices and firm profits, which further amplifies the dynamics. The presence of positive feedback from assets prices to profits depends on the relative strength of general equilibrium forces, which in the model in this section operate through the real wage.

I now turn to the development of a richer model that embeds the same mechanism, but can also be taken to the data to study its quantitative implications.

2.5 FULL MODEL FOR QUANTITATIVE ANALYSIS

This section embeds the mechanism discussed previously into a New-Keynesian business cycle model with a financial accelerator. Compared to the simple model in the previous section, there are a number of new elements. First, capital no longer has to be financed entirely out of debt. Instead, I allow for endogenous fluctuations in net worth. To prevent firms from saving until they become unconstrained, I impose exogenous entry and exit, as in Bernanke, Gertler and Gilchrist (1999). Second, I provide a microfoundation of the borrowing constraint by means of a limited commitment problem. Third, I add several standard business cycle frictions: nominal rigidities, which enables me to introduce monetary policy and later on analyse its effects on welfare under learning; and investment adjustment costs, which allow for a better fit of the model.

2.5.1 *Model setup*

The economy is closed and operates in discrete time. There are a number of different agents:

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

2. *Firm owners* only consume differentiated final goods. They trade shares in intermediate goods producers and receive dividend payments.

3. *Households* consume differentiated final goods and supply homogeneous labour to labour agencies. They lend funds to intermediate goods producers.

4. *Labour agencies* transform homogeneous household labour into differentiated labour services, which they sell to intermediate goods producers. They are owned by households.

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

6. *Capital goods producers* produce new capital goods from final consumption goods subject to an investment adjustment cost.

7. *The fiscal authority* sets certain tax rates to offset steady-state distortions from monopolistic competition.

8. *The central bank* sets nominal interest rates.

Since most elements of the model are standard, I focus on the financially constrained firms, firm owners, households and the central bank. Additional details are provided in Appendix A.2.

2.5.1.1 *Households*

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

$$\max_{\left(C_t, L_t, B_{jt}, B_t^g\right)_{t=0}^{\infty}} \mathbb{E}_0^{\mathcal{P}} \sum_{t=0}^{\infty} \beta^t u\left(C_t, L_t\right)$$

$$\begin{aligned} \text{s.t. } C_t &= \tilde{w}_t L_t + B_t^g - \left(1 + i_{t-1}\right) \frac{p_{t-1}}{p_t} B_{t-1}^g \\ &+ \int_0^1 \left(B_{jt} - R_{jt-1} B_{jt-1}\right) dj + \Pi_t \end{aligned} \tag{2.37}$$

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

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

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

$$p_t = \left(\int_0^1 (p_{it})^{1-\sigma} \, di \right)^{\frac{1}{1-\sigma}} ; \, C_{it} = \left(\frac{p_{it}}{p_t} \right)^{-\sigma} C_t. \tag{2.38}$$

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

$$\tilde{w}_t = -\frac{u_{Lt}}{u_{Ct}} \tag{2.39}$$

$$1 = \mathbb{E}_t^{\mathcal{P}} \beta \frac{u_{Ct+1}}{u_{Ct}} \frac{1+i_t}{\pi_t}. \tag{2.40}$$

We can define the stochastic discount factor of the households as $\Lambda_t = \beta u_{Ct+1} / u_{Ct}$.

2.5.1.2 *Central bank*

Like most of the New-Keynesian literature, the model is cashless, with the central bank affecting allocations in the presence of nominal rigidities by setting the nominal interest rate. In the baseline version of the model, I assume that the central bank conducts monetary policy through the use of a Taylor-type interest rate rule:

$$i_t = \rho_i i_{t-1} + (1 - \rho_i) \left(1/\beta + \pi_t + \phi_\pi \left(\pi_t - \pi_t^* \right) \right) \tag{2.41}$$

where π_t^* is the central bank's (time-varying) inflation target, ρ_i is the degree of interest rate smoothing and $\phi_\pi > 1$.

2.5.1.3 *Intermediate good producers (firms)*

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

$$Y_{jt} = \left(K_{jt-1} \right)^\alpha \left(A_t L_{jt} \right)^{1-\alpha}, \tag{2.42}$$

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

$$L_{jt} = \max_{L_{jht}} \left(\int_0^1 \left(L_{jht} \right)^{\frac{\sigma_w - 1}{\sigma_w}} dh \right)^{\frac{\sigma_w}{\sigma_w - 1}} \tag{2.43}$$

$$\text{s.t. } w_t p_t L_{jt} = \int_0^1 W_{jht} L_{jht} dh \tag{2.44}$$

The firm's problem can then be treated as if the labour index was acquired in a competitive market at the real wage index w_t.[15] Output is sold competitively to final good producers at price q_t. During production, the capital stock depreciates at rate δ. This depreciated capital can be traded by the firm at the price Q_t.

At this point, the net worth of the firm is the difference between the value of its assets and its outstanding debt:

$$N_{jt} = q_t Y_{jt} - w_t L_{jt} + Q_t (1-\delta) K_{jt-1} - R_{jt-1} B_{jt-1} \tag{2.45}$$

I assume that the firm exits with a 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 ζ of its earnings as dividends, where earnings are given by $E_{jt} = N_{jt} - Q_t K_{jt-1} + B_{jt-1}$.[16] If it exits, it must pay out its entire net worth as dividends. It is subsequently replaced by a new firm which receives the index j. I assume that this new firm gets endowed with a fixed number of shares, normalised to one, and is able to raise an initial amount of net worth. This amount equals $\omega (N_t - \zeta E_t)$ where $\omega \in (0,1)$ and N_t and E_t are aggregate net worth and earnings, respectively.

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

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

This firm then decides on the new stock of debt B_{jt} and the new capital stock K_{jt}. Its balance sheet must satisfy:

$$Q_t K_{jt} = B_t^j + \tilde{N}_{tj} \tag{2.46}$$

where the price of capital Q_t can vary in the presence of adjustment costs.

Firms maximise the present discounted value of their dividend payments using the discount factor of their owners. In doing so, they face financial constraints. Before describing these constraints though, I first turn to the description of the firms' owners.

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

16 The optimal dividend policy in this model would be to never pay dividends until exit. In this case, aggregate dividends would be proportional to aggregate net worth. This implies a dividend process that is not nearly as volatile as in the data, and thus makes it impossible to obtain good asset pricing properties even under learning. Imposing that firms need to pay out a fraction of their earnings greatly improves the quantitative fit of the model.

2.5.1.4 *Firm owners*

Firm owners differ from households in their capacity to own intermediate firms. The representative firm owner is risk-neutral and discounts future income at the rate $\tilde{\beta} = \beta G^{-\theta} < \beta$, with G being the growth rate of consumption in the non-stochastic steady state. He can buy shares in firms indexed by $j \in [0,1]$. As described above, when a firm exits it pays out its net worth N_{jt} as dividends, and is replaced by a new firm which raises equity $\omega\left(N_t - \zeta E_t\right)$. Let the set of exiting firms in each period t be denoted by $\Gamma_t \subset [0,1]$. Then, the firm owner's utility maximisation problem is given by:

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

$$\begin{aligned} \text{s.t. } C_t^f &= -\int_0^1 S_{jt} P_{jt} dj + \int_{j \notin \Gamma_t} S_{jt-1}\left(P_{jt} + D_{jt}\right) dj \\ &+ \int_{j \in \Gamma_t} \left[S_{jt-1} D_{jt} - \omega\left(N_t - \zeta E_t\right) + P_{jt}\right] dj \qquad (2.47) \\ S_t^j &\in [0, \bar{S}] \qquad (2.48) \end{aligned}$$

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

The first term on the right hand side of the budget constraint deals with continuing firms and is standard: Each share in firm j pays dividends D_{jt} and continues to trade, at price P_{jt}. The second term deals with firm entry and exit. If the household owns a share in the exiting firm j, he receives a terminal dividend. The firm is then delisted in the stock market, and so $S_{jt-1}P_{jt}$ does not appear. At the same time, a new firm j appears which is able to raise a limited amount of equity $\omega\left(N_t - \zeta E_t\right)$ from the firm owner in exchange for a unit amount of shares that can be traded at price P_{jt}. In addition, upper and lower bounds on traded stock holdings are introduced to make firm owners' demand for stocks finite under arbitrary beliefs, as in the stylised model of the previous section. In equilibrium, they are never binding.

The first order conditions of the firm owner are

$$S_{jt} \begin{cases} = 0 & \text{if } P_{jt} > \tilde{\beta} \mathbb{E}_t^{\mathcal{P}}\left[D_{jt+1} + P_{jt+1} \mathbb{1}_{\{j \notin \Gamma_{t+1}\}}\right] \\ \in [0, \bar{S}] & \text{if } P_{jt} = \tilde{\beta} \mathbb{E}_t^{\mathcal{P}}\left[D_{jt+1} + P_{jt+1} \mathbb{1}_{\{j \notin \Gamma_{t+1}\}}\right] \\ = \bar{S} & \text{if } P_{jt} < \tilde{\beta} \mathbb{E}_t^{\mathcal{P}}\left[D_{jt+1} + P_{jt+1} \mathbb{1}_{\{j \notin \Gamma_{t+1}\}}\right] \end{cases} \qquad (2.49)$$

2.5.1.5 *Borrowing constraint*

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

Each period, lenders (households) and borrowers (firms) meet to decide on the lending of funds. Pairings are anonymous to rule out repeated interactions. The incompleteness of contracts imposed is that 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 realisation 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 the seizure of the firm by the lender, in which case the current firm owners receive zero.

The lender, a household, does not have the ability to run the firm though. The usual assumption in the literature is that she has to liquidate the firm's asset in this case. In this model, the lender can always liquidate as well. In this case, all debt and a fraction $1 - \xi$ of the firm's capital is destroyed. The remaining capital can be sold in the next period, resulting in a total recovery value of $\xi Q_{t+1} K_{jt}$. On top of this, with some probability x (independent across time and firms), the lender gets the opportunity to "restructure" the firm. 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. It will turn out in equilibrium that the recovery value in this case is just $\xi \left(P_{jt} + B_{jt}\right)$ and that lenders always prefer restructuring to liquidation. The debt contract then takes the form of a leverage constraint in which total firm value is a weighted average of liquidation and market value:

$$B_{jt} \leq \xi \left(x \underbrace{\mathbb{E}_t^{\mathcal{P}} \Lambda_{t+1} Q_{t+1} \xi K_{jt}}_{\text{liquidation value}} + (1 - x) \underbrace{\left(P_{jt} + B_{jt}\right)}_{\text{market value}} \right) \tag{2.50}$$

2.5.1.6 *Further model elements and market clearing*

Final good producers, indexed by $i \in [0, 1]$, combine the homogeneous intermediate good into a differentiated final good using a one-for-one technology. Their revenue is subsidised by the government at the rate τ.[17] Per-period profits of producer i are given by $\Pi_{it}^{Y} =$

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

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

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

Finally, the exogenous stochastic processes are productivity and the inflation target shock:

$$\log A_t = (1-\rho)\log \bar{A} + \rho \log A_{t-1} + \log \varepsilon_{At} \tag{2.51}$$

$$\pi^*_t = \rho_\pi \pi^*_{t-1} + \log \varepsilon_{\pi t} \tag{2.52}$$

$$\varepsilon_{At} \sim \mathcal{N}\left(0, \sigma^2_A\right) \tag{2.53}$$

$$\varepsilon_{\pi t} \sim \mathcal{N}\left(0, \sigma^2_\pi\right) \tag{2.54}$$

Market clearing needs to take into account the distortions from price and wage dispersion. All market clearing conditions are listed in Appendix A.2.

2.5.2 *Rational expectations equilibrium*

I first describe the equilibrium under rational expectations. An equilibrium is a set of stochastic processes for prices and allocations, a set of strategies in the limited commitment game, and an expectation measure $\mathcal{P}$ such that the following holds for all states and time periods: Markets clear; allocations solve the optimisation programmes 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* probability measure induced by the equilibrium.

Under appropriate parameter restrictions, there exists a rational expectations equilibrium characterised by the following properties (proofs and characterisation of the restrictions are relegated to Appendix A.3):

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

$$\begin{aligned} Y_t &= \alpha K_{t-1}^{\alpha} \left(A_t \tilde{L}_t\right)^{1-\alpha} && (2.55) \\ R_t^k &= q_t \alpha \frac{Y_t}{K_{t-1}} + Q_t (1-\delta) K_{t-1} && (2.56) \end{aligned}$$

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

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

$$P_t = \int_0^1 P_{jt} = \tilde{\beta} \mathbb{E}_t \left[D_{t+1} + \frac{1-\gamma}{1-\gamma+\gamma\omega} P_{t+1} \right]. \qquad (2.57)$$

The "correction" term in the continuation value P_{t+1} can be understood as follows. The stock market index P_t is the value of all currently existing firms, not including firms that are not yet born. In $t+1$, a fraction γ of firms exit and pay dividends γN_t. The remaining firms are left with net worth $(1-\gamma) N_{t+1}$ and pay dividends $(1-\gamma)\zeta E_t$ and , but a mass γ of new firms also enters, each endowed with initial net worth $\omega (N_t - \zeta E_t)$.

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

$$R_{jt} = R_t = \left(\mathbb{E}_t \beta u_{Ct+1}/u_{Ct}\right)^{-1}. \qquad (2.58)$$

The lender only accepts debt payments up to a certain limit $\bar{B}_{jt}$. The firm always exhausts this limit, $B_{jt} = \bar{B}_{jt}$, which is proportional to the firm's net worth $\tilde{N}_{jt}$. If the firm defaulted and the lender seized the firm, she would always prefer restructuring to liquidation. Intuitively, this is because capital is more valuable inside the firm than outside of it: Because acquiring capital is difficult due to financial frictions, firm owners will always pay the lender a higher price for a restructured, operational firm than for its capital stock alone.

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 and evolve as

$$\begin{aligned} N_t &= R_t^k K_{t-1} - R_{t-1} B_{t-1} && (2.59) \\ Q_t K_t &= (1-\gamma+\gamma\omega)\left((1-\zeta) N_t + \zeta \left(B_{t-1} - Q_t K_{t-1}\right)\right) && \\ &+ B_t && (2.60) \\ B_t &= x \mathbb{E}_t \Lambda_{t+1} Q_{t+1} \xi K_t + (1-x)\, \xi \left(P_t + B_t\right). && (2.61) \end{aligned}$$

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

2.5.3 *Learning equilibrium*

I introduce learning about stock market valuations as in the simple model of Section 2.4. One slight complication is now that there is a continuum of firms to be priced in the market. In this respect, I retain the belief that the stock price of an individual firm is proportional to firm net worth, as is the case under rational expectations. As such, under $\mathcal{P}$,

$$P_{jt} = \frac{N_{jt}}{N_t} P_t. \tag{2.62}$$

But while investors know how to price individual stocks by observing the valuation of the market, they are uncertain about the evolution of the market itself. As in the simple model of the previous section, I impose the same beliefs about aggregate stock prices as in the last section along with the other assumptions (equations (2.25)-(2.28)), including conditionally model-consistent expectations: For any variable x_t other than the aggregate stock price, any future date $t+\tau$, and any sequence $P_{t+1}, \ldots, P_{t+\tau}$ which is on the equilibrium path, agents' beliefs coincide with the best statistical prediction of x_t *conditional* on the realisation of stock prices: $\mathbb{E}_t^{\mathcal{P}}\left[x_{t+\tau} \mid P_{t+1}, \ldots, P_{t+\tau}\right] = \mathbb{E}_t\left[x_{t+\tau} \mid P_t, P_{t+1}, \ldots, P_{t+\tau}\right]$.

In practice, I solve the model using the two-stage procedure described in Appendix A.1.

Some readers might object that the belief system about stock prices is misspecified in the sense that agents are never able to learn the true law of motion for P_t (where P_t is not a random walk and depends on several state variables). However, this is a deliberate choice. The chosen form of subjective beliefs captures several aspects of reality: People think that stock prices are well approximated by random walks; still, they try to identify predictability in prices; their subjective expectations seem to ignore the degree of mean reversion in returns observed in the data, instead overly extrapolating past observation; and they think that the laws governing prices are changing over time, so that observations from the distant past are only of limited usefulness in predicting the future.

2.5.4 *Frictionless benchmark*

Because my aim is to gauge the importance of financial frictions for business cycle analysis, I will also use a benchmark model without financial frictions to which the model can be compared to. This benchmark model will be identical to the rational expectations model above, except that intermediate firms are now owned by households, and face no frictions in accessing external finance.

2.5.5 *Choice of parameters*

I partition the set of parameters into two groups. The first set of parameters is calibrated to first moments, and the second set is estimated by simulated method of moments.

2.5.5.1 *Calibration*

The capital share in production is set to $\alpha = .33$, implying a labour share in output of two thirds. The depreciation rate $\delta = .025$ corresponds to 10% annual depreciation. The non-stochastic trend productivity growth rate G is set to its post-war average of 1.64% annually. The persistence of the temporary component of productivity is set to .95.

The discount factor of the household is set such that the steady-state interest rate matches the average annual real return on Treasury bills of 2%, implying a discount factor $\beta = .9991$. The firm owners' discount factor is set such that in the non-stochastic steady state, bonds and stocks have the same return ($\tilde{\beta} = .995$). The elasticity of substitution between varieties of the final consumption good, as well as that between varieties of labour used in production, is set to $\sigma = \sigma_w = 4$. The Frisch elasticity of labour supply is set to three, implying $\phi = .33$.

The strength of monetary policy reaction to inflation is set to $\phi_\pi = 1.5$, while the degree of nominal rate smoothing is set to $\rho_i = 0.5$. The inflation target is set to $\pi^* = 0$. The persistence of the inflation target is set to $\rho_\pi = .99$.

Four parameters describe the structure of financial constraints: x, the probability of restructuring after default; ξ, the tightness of the borrowing constraint; ω, the equity received by new firms relative to average equity; and γ, the rate of firm exit and entry. I calibrate the restructuring rate x to equal 9.3%. This is the fraction of US business bankruptcy filings in 2006 which filed for Chapter 11 instead of Chapter 7, and which subsequently emerged from bankruptcy with an approved restructuring plan. I have to restrict myself to 2006 because it is the only year for which this number can be constructed. Public bankruptcy data shows that the fraction of Chapter 11 cases as opposed to Chapter 7 cases fluctuates around 28%, and several papers analysing sub-samples of Chapter 11 filings arrive at confirmation rates between 29% and 64%, which suggests that the 2006 number is reasonable (a sensitivity check is included in Section 2.6.5).[18] The

18 Data on bankruptcies by chapter are available at http://www.uscourts.gov/Statistics/BankruptcyStatistics.aspx. Data on Chapter 11 outcomes are ana-

remaining three parameters are chosen such that the non-stochastic steady state of the model jointly matches the US average investment share in output of 20%, average debt-to-equity ratio of 1:1 (as recorded in the Fed Flow of Funds), and average quarterly P/D ratio of 139 (taken from the S&P500). The parameter values thus are $\gamma = 0.010$, $\xi = 0.38$, and $\omega = 0.55$. Intuitively, a larger value of ξ relaxes the borrowing constraint, leading to higher leverage. A larger value of ω means that new firms enter with more equity, increasing aggregate net worth and capital. Since returns on capital are diminishing, this raises the investment share in output. A higher rate of exit γ makes firms shorter-lived, reducing the firm's value relative to its current dividend payments, thereby lowering the P/D ratio.

2.5.5.2 *Estimation*

The remaining parameters are the standard deviations of the technology and monetary shocks (σ_A, σ_π), the size of investment adjustment costs (ψ), the degree of nominal price and wage rigidities (κ, κ_w), the fraction of dividends paid out as earnings by continuing firms (ζ), and the learning gain (g). Since my goal is to see how well the model can do in terms of matching both business cycle and asset pricing facts, I estimate these parameters to minimise the distance to a set of second moments pertaining to both. I use the variances of the nominal interest rate, the inflation rate, real output, consumption, investment, employment, as well as real dividend payments and stock prices in the data.[19] The set of estimated parameters θ solves

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

where $m\left(\theta\right)$ are moments obtained from model simulation paths with 50,000 periods, $\hat{m}$ are the estimated moments in the data, and W is a weighting matrix.[20] I also impose that θ has to lie in a subset $\mathcal{A}$ of the parameter space which rules out deterministic oscillations of impulse response functions.[21] Parameters outside this region would fit the moments well but can be ruled out from prior knowledge about

lysed in various samples by Flynn and Crewson (2009), Warren and Westbrook (2009), Lawton (2012), and Altman (2014).

19 All variables are at quarterly frequency and HP-detrended. I use CPI inflation as the inflation measure and the Federal Funds rate as the nominal interest rate. Employment is total non-farm payroll employment. Consumption is the sum of services and non-durable private consumption. Investment is the sum of private non-residential fixed investment and durable consumption. Output is the sum of consumption and investment. Dividend payments are the four-quarter moving average of S&P 500 dividends and stock prices is the S&P500 index.

20 I choose $W = \text{diag}\left(\hat{\Sigma}\right)^{-1}$ where $\hat{\Sigma}$ is the covariance matrix of the data moments, estimated using a Newey-West kernel with optimal lag order. This choice of W leads to a consistent (albeit not fully efficient) estimator that places more weight on moments which are more precisely estimated in the data.

21 To be precise, $\theta \notin \mathcal{A}$ 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 2.3: Estimated parameters.

param.	σ_a	σ_{π^*}	ψ	κ	κ_w	ζ	g
learning	.00705	$8.53 \cdot 10^{-4}$	8.51	.776	.946	.765	.0046
	(.00157)	$(2.94 \cdot 10^{-4})$	(2.05)	(.165)	(.025)	(.027)	$(8.66 \cdot 10^{-4})$
RE	.0114	.00117	.168	.914	0	.637	-
	(.00201)	$(2.3 \cdot 10^{-4})$	(.715)	(.0216)	(4.45)	(.309)	-

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

the shape of impulse responses. Table 2.3 summarises the results of the SMM estimation. The first row presents the results under learning. The SMM procedure selects a rather high degree of adjustment costs (although imprecisely estimated) and of nominal rigidities. The fraction of earnings paid out as dividends is fitted to more than 80%, higher than the actual historical average for the S&P500 at about 50%. This suggests that my assumption about the behaviour of dividends is overly simplistic. This is a commonly encountered problem (Covas and Den Haan, 2012) which I have to leave for future research. Finally, a low estimate for the learning gain g implies that agents believe the predictability of stock prices to be small.

The second column presents the parameters as estimated under rational expectations (therefore not including the learning gain g). A much smaller degree of investment adjustment costs and nominal rigidities is needed here to fit the model. Nevertheless, the size of the standard errors of the two shocks are more than 20% larger than under learning. This indicates greatly increased amplification of business cycle shocks in the model with learning.

2.6 RESULTS

2.6.1 *Business cycle and asset price moments*

To get a better understanding of the quantitative properties of the model, Table 2.4 reviews key business cycle moments as well as the statistics describing stock price volatility and return predictability across model specifications. The moments for the estimated learning model are in Column (1), while Columns (2) and (3) contain the corresponding moments for the model under rational expectations and the frictionless benchmark. Here, the parameters are held constant at the same values as for the learning model. By nature of the

Table 2.4: Comparing moments in the data and across model specifications.

	moment	data	(1) learning	(2) RE	(3) fric.less	(4) RE re-est.
output volatility	$\sigma_{hp}(Y_t)$	1.43%	1.41%	.80%	.55%	1.37%
relative volatility	$\sigma_{hp}(C_t)/\sigma_{hp}(Y_t)$	0.60	.63	1.09	1.43	.66
	$\sigma_{hp}(I_t)/\sigma_{hp}(Y_t)$	2.90	3.10	.65	.34	3.16
	$\sigma_{hp}(L_t)/\sigma_{hp}(Y_t)$	1.13	1.16	.98	1.02	1.24
	$\sigma_{hp}(\pi_t)/\sigma_{hp}(Y_t)$	.19	.22	.34	.47	.23
	$\sigma_{hp}(i_t)/\sigma_{hp}(Y_t)$	.20	.20	.32	.43	.20
	$\sigma_{hp}(D_t)/\sigma_{hp}(Y_t)$	2.99	3.07	2.36	-	1.11
correlation with output	$\rho_{hp}(C_t, Y_t)$	.94	.92	.99	.99	.94
	$\rho_{hp}(I_t, Y_t)$	.95	.93	.99	.99	.94
	$\rho_{hp}(L_t, Y_t)$	.88	.91	.64	.30	.92
excess volatility	$\sigma_{hp}(P_t)/\sigma_{hp}(D_t)$	2.63	1.81	.21	-	.15
	$\sigma\left(\frac{P_t}{D_t}\right)$	.408	.131	.037	-	.069
	$\sigma\left(R^{stock}_{t,t+1}\right)$	.335	.150	.011	-	.019
return predictability	$\rho\left(\frac{P_t}{D_t}, R^{stock}_{t,t+4}\right)$	-.218	-.240	-.002	-	.002
	$\rho\left(\frac{P_t}{D_t}, R^{stock}_{t,t+20}\right)$	-.439	-.646	.015	-	.027
	$\rho\left(\frac{P_t}{D_t}, \frac{P_{t+4}}{D_{t+4}}\right)$	.907	.432	.758	-	.681

Quarterly US data 1962Q1-2012Q4. π_t is quarterly CPI inflation. i_t is the Federal Funds rate. L_t is total non-farm payroll employment. Consumption C_t consists of 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. The stock price index P_t is the S&P500.

estimation, the learning model has the best fit across Columns (1) to (3). The comparison serves to single out the contribution of learning and financial frictions to the fit. Column (4) presents the moments from the re-estimated model under rational expectations.

The first row reports the standard deviation of detrended output. By this measure, output fluctuations under learning are double the size of those under rational expectations; in other words, learning adds considerable endogenous amplification to the model. In fact, the standard deviations of shocks in Column (4) needed to generate the same amount of output volatility are much larger.

The following rows report the relative volatilities of key macroeconomic time series. The learning model matches all of them well. Shutting down learning leads to a sharp drop in the volatility of investment and a corresponding increase in the volatility of consumption. This is because the estimated learning model features a rather high level of investment adjustment costs to match investment volatility. Without high adjustment costs, the learning model would transform the high degree of observed asset price volatility into a counterfactually high degree of investment volatility, i.e. the amplification mechanism would be too strong. When the amplification is shut down in Columns (2) and (3), investment becomes very smooth and as a flip side, consumption is prevented from being smoothed. Looking at inflation and interest rates, shutting down learning increases their volatility relative to output, although this is mostly due to the decrease in output volatility itself. The volatility of employment and dividends is not much affected. Finally, the re-estimated model in Column (4) is able to match all moments just as well as under learning with the exception of dividends.

Big differences appear in the asset price statistics. The learning model is able to approach the volatility of prices and returns to a degree which is impossible to achieve even in the re-estimated RE model.[22] The ability to generate realistic stock price volatility is of course at the heart of the amplification mechanism, since stock market valuations enter firms' borrowing constraints. Moreover, returns are negatively predicted by the P/D ratio. At the one year horizon, the predictability is very similar to that found in the data. At the five-year horizon however, it is too high. This is also reflected in the fact that the P/D ratio decays too fast, as documented in the last row. In a sense, this is not really surprising: The learning model has only one parameter (the learning gain g) to match all statistics pertaining to stock prices. In particular, it can be shown that a higher gain increases volatility, but reduces persistence of prices. A richer belief specification could possibly achieve a better fit, but at the expense of introducing additional parameters.

2.6.2 *Impulse response functions*

Looking at impulse response functions reveals some of the workings of the amplification mechanism at play. Figure 2.6 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 frictionless benchmark. Looking at output Y_t, one can see quite clearly how the rational expectations version generates only a slightly amplified re-

22 In fact, it would be possible to obtain an even better fit of the asset pricing moments, but at the expense of business cycle moments.

Figure 2.6: Impulse responses to a persistent productivity shock.

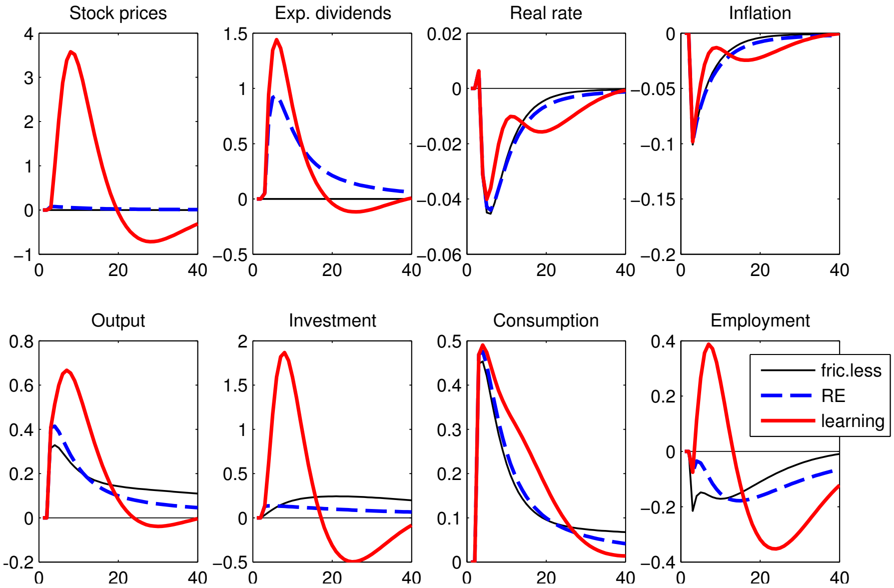

Impulse responses to a one-standard deviation innovation in ε_{At}. Stock prices P_t, dividends D_t, output Y_t, investment I_t, consumption C_t, and employment L_t are in 100*log deviations. The interest rate i_t and inflation π_t are in percentage-point deviations.

sponse to the shock compared to the frictionless benchmark. With learning, however, output rises by almost double the amount and the response is hump-shaped. Consumption C_t, and investment I_t also exhibit much stronger responses, and employment L_t also rises considerably (without learning it contracts). This amplification is due to the large rise in the stock prices P_t, even though the degree of dependence of borrowing constraints on stock prices (measured by x) is less than 10%. Under rational expectations, stock prices stay virtually flat, thus producing almost no amplification through the equity price channel. The behaviour of dividends D_t reveals that initially dividends fall, partly offsetting the rise of beliefs, but then overshoot their rational expectations counterpart.

The feedback loop is also present after an interest rate shock. Figure 2.7 plots the response to a temporary reduction in the nominal interest rate. Again, all macroeconomic aggregates rise substantially more under learning than under both rational expectations and the frictionless benchmark. The monetary stimulus increases stock prices and thus relaxes credit constraints. The consequent increase in investment demand further raises immediate inflationary pressure, which in the model leads to the central bank undoing its interest rate reaction, acting according to its interest rate rule. Here, asset prices undershoot after the stimulus and remain depressed for a long period of time, leading to lower output, inflation and employment with a trough at about 16 quarters after the stimulus. This boom and bust pattern is not present under rational expectations.

Figure 2.7: Impulse responses to a monetary shock.

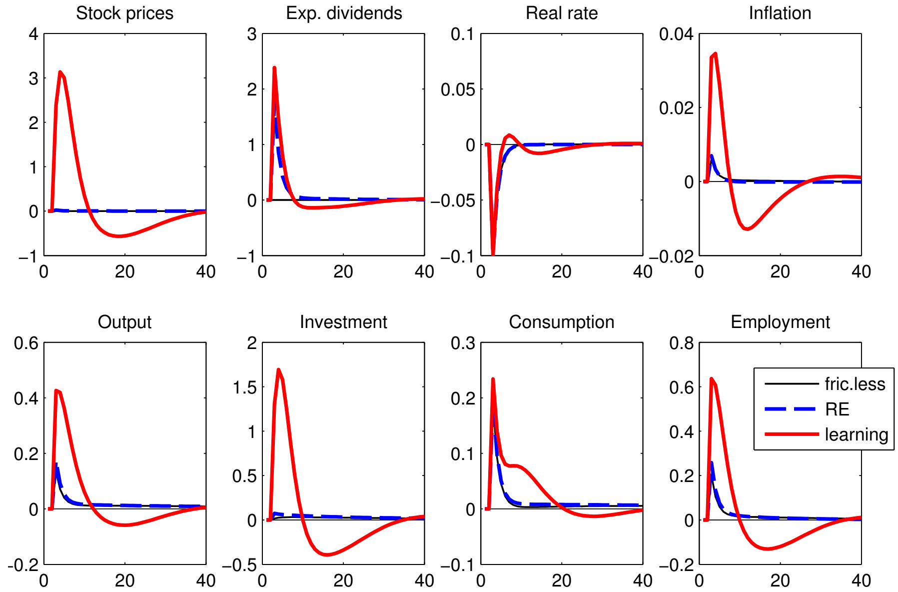

Impulse responses to a one-standard deviation innovation in ε_{mt}. Stock prices V_t, dividends D_t, output Y_t, investment I_t, consumption C_t and employment L_t are in 100*log deviations. The interest rate i_t and inflation pi_t are in percentage point deviations.

2.6.3 *Does learning matter?*

The discussion so far has mainly focused on how large swings in asset prices lead to large swings in real activity through their effect on credit constraints. This raises the question of whether learning is necessary for the story at all - maybe any theory which replicates the same asset price dynamics also replicates the same business cycle outcomes? To answer it, I set up the following experiment. I replace the stock market value P_t in the borrowing constraint (2.61) with an exogenous process V_t that has the same law of motion as the stock price under learning. More precisely, I fit an ARMA(10,5)-process for V_t such that its impulse responses are as close as possible to those of P_t under learning (the exogenous shock in the ARMA-process are the productivity and monetary shocks). I then solve this model, but with rational expectations. If learning only matters because it affects stock price dynamics, then this hypothetical model should have identical dynamics to the model under learning.

Figure 2.8 shows that this is not the case. The ARMA-process fits stock prices well: The impulse response of P_t under learning and V_t in the counterfactual experiment are indistinguishable. But after a positive productivity shock, output, investment and consumption rise much more under learning, even though the counterfactual model has the same stock price dynamics by construction. The reason is found in the fact that expectations matter beyond stock prices: For interest rates equilibrating loan demand by firms and supply by house-

Figure 2.8: Does learning matter?

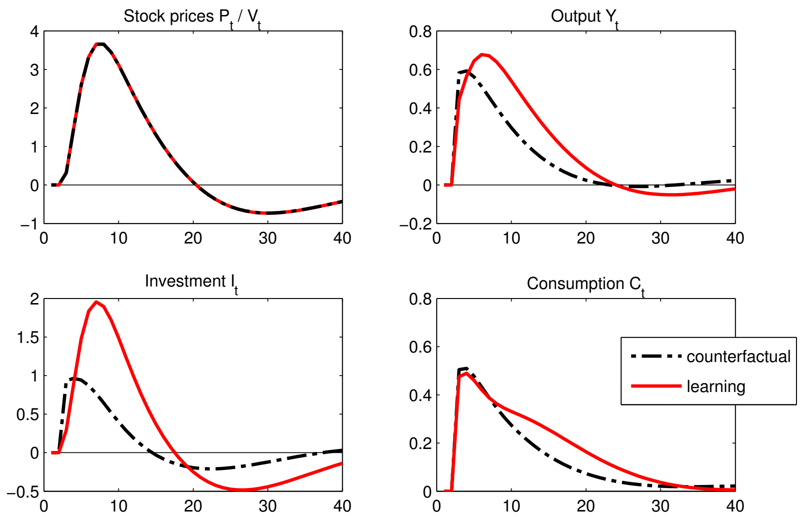

Solid red line: Impulse response to a one-standard deviation positive productivity shock under learning. Black dash-dotted line: Impulse response to a hypothetical rational expectations model with stock price dynamics identical to those under learning (see text).

holds; for inflation and wages, set by forward-looking Calvo price and wage setters; and for the borrowing constraint (2.61) itself, since it depends on the expected liquidation value of capital $\mathbb{E}_t Q_{t+1}$. This last channel is in fact crucial for the additional amplification. Under learning, agents do not understand that the increase in stock prices is temporary. By over-predicting the slackness of borrowing constraints, they also over-predict the amount of investment taking place in the future, and hence the future price of capital. Lenders, predicting a high value of the firm's capital stock, are then willing to lend more to firms today. Thus, identical firm market value still leads to easier access to credit under learning.

This illustrates how expectations in financial markets, over and above their effect on asset prices, can have important effects on the real economy.

2.6.4 *Relation to survey evidence on expectations*

Agents in the learning equilibrium make systematic, predictable forecast errors. The patterns of predictability are testable model implications. As it turns out, they are surprisingly consistent with survey data.

Importantly, agents in the model do make systematic forecast errors not only about stock prices, but also about almost all other endogenous model variables. This is despite the fact that, *conditional* on stock prices, agents' beliefs are model-consistent. A systematic

Figure 2.9: Return expectations and expected returns in a model simulation.

Expected and realised nominal returns along a simulated path of model with learning. Simulation length 200 periods. Theoretical correlation coefficient for expected returns $\rho = .30$, for realised returns $\rho = -.11$.

mistake in predicting stock prices will still translate into a corresponding mistake in predicting the tightness of borrowing constraints, and hence investment, output, and so forth. I can therefore compare forecast errors on many model variables with the data. Predictability in all variables arises from the introduction of only a single parameter (the learning gain), so this is a potentially tough test for the model.

Figure 2.9 repeats the scatter plot of Section 2.3, contrasting expected and realised one-year ahead returns in a model simulation. The same pattern as in the data emerges: When the P/D ratio is high, return expectations are most optimistic. In the learning model, this has a causal interpretation: high return expectations drive up stock prices. At the same time, realised future returns are on average low when the P/D ratio is high. This is because the P/D ratio is mean-reverting (which agents do not realise, instead extrapolating past price growth into the future): At the peak of investor optimism, realised price growth is already reversing and expectations are due to be revised downward, pushing down prices towards their long-run mean.

Table 2.5 compares the correlation of forecast errors with the predictors discussed in Section 2.3, in the data and the model with learning. Note that the correlations under rational expectations need not be reported, as they are all zero by construction.

Column (1) reports predictability based on the P/D ratio. As already discussed, the correlation is negative, and it is even stronger in the data than in the model. The correlation is also negative for the other macroeconomic aggregates, again with a higher magnitude in the data. Column (2) repeats the exercise for the growth rate of the P/D ratio. This measure positively predicts forecast errors: When the stock market is rising, people are on average not optimistic enough about

Table 2.5: Forecast errors under learning and in the data.

	(1)		(2)		(3)	
forecast variable	$\log PD_t$		$\Delta \log PD_t$		forecast revision	
	model	data	model	data	model	data
$R^{stock}_{t,t+4}$	-.22	-.44	.30	.06	-.19	n/a
$Y_{t,t+3}$	-.10	-.21	.22	.22	.21	.29
$I_{t,t+3}$	-.14	-.20	.36	.25	.36	.31
$C_{t,t+3}$	-.04	-.19	.02	.21	.02	.23
$u_{t,t+3}$	.10	.05	-.25	-.27	.15	.43

Data as in Table 2.2. Model correlations are exact theoretical correlations. Stock returns are nominal returns. Unemployment in the model is taken to be $u_t = 1 - L_t$.

returns and economic activity. This is because expectations about asset prices (and hence lending conditions) adjust only slowly. Here, the degree of correspondence with the data is striking. Only forecast errors on consumption remain too weakly predictable in the model.

The model also does very well in terms of the Coibion and Gorodnichenko (2010) predictability statistic. As documented in Column (3), forecast errors on macro aggregates are predictable by the direction of the forecast revision, again reflecting slow belief adjustment. This is also true in the model. For stock returns, the model predicts a negative correlation of return forecast errors with the revision, but the CFO survey does not allow for the construction of a corresponding statistic.

Forecast error predictability is illustrated graphically in Figure 2.10. The solid red line is a standard impulse response function to a tech-

Figure 2.10: Actual versus expected impulse response.

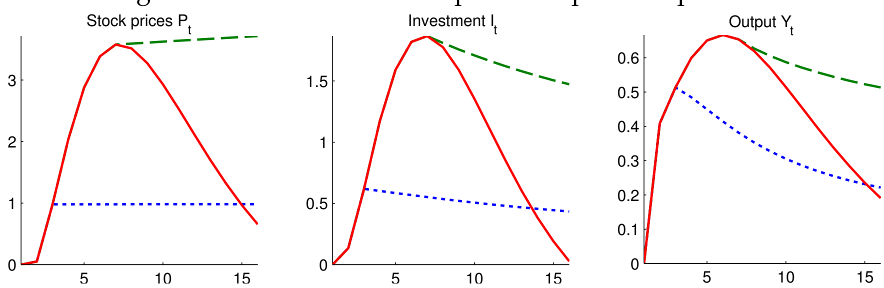

Solid red line: Impulse response to a one-standard deviation positive productivity shock in Period 1. Green dashed line: mean subjective forecasts taken in Period 7. Blue dotted line: mean subjective forecasts taken in Period 3.

nology shock in the learning model. Consider the following thought experiment. Suppose that at the time of the shock, the economy is in steady state and that no further shocks are realised. Under rational expectations, the mean expectation about the future path of the economy at any point in time is then equal to the impulse response itself. Under learning however, subjective expectations do not coincide with the impulse response. The green dashed line of Figure 2.5 depicts the mean subjective forecast at the peak of the stock market. Agents do not foresee the decline in the stock market, and instead extrapolate high stock price growth into the future. Because stock market valuations matter for access to credit, agents also forecast loose borrowing conditions and are too optimistic about investment and output as well.[23] The green dashed line (expectations) is above the red solid line (realisation) when the P/D ratio is high: It negatively predicts forecast errors. Next, the change in the P/D ratio predicts forecast errors the other way around, which is illustrated with the blue dotted line. This is the forecast made at a time in which the P/D ratio is rising fast. In this situation, agents under-predict the size of the coming boom in the stock market and real activity. The blue dotted line is below the red solid line: Forecast errors are positive when P/D ratio growth is positive, too.

2.6.5 *Sensitivity checks*

The amount of endogenous amplification under learning relies on the dynamics of stock prices, and also on general equilibrium effects (as was already illustrated in the simplified model). To gain an idea of the sensitivity of the results, Figure 2.11 plots the volatility of output and stock prices as a function of a set of influential parameters.

Panel (a) shows the role of the restructuring rate x which parametrises the dependency of borrowing constraints on firm market value (as opposed to the liquidation value of its capital). Not surprisingly, this parameter is key in driving amplification under learning. The point $x = 0$ is a special case. At this point, stock prices have no allocative role for the economy. The model dynamics under learning and rational expectations then coincide perfectly. Another special point is $x = 1$, where borrowing constraints depend exclusively on stock prices. This was the case analysed in the simple model of Section 2.4. In the full model, this degree of stock price dependency leads to so much amplification that the belief dynamics become explosive, and no stable equilibrium exists.[24] Another remarkable fact is that

23 Note that the forecasts for output and investment are still downward-sloping as agents are perfectly aware of the mean reversion of productivity. Their long-run forecast has permanently higher stock prices, output and investment because of easier access to credit, while productivity remains at the steady-state level.

24 To be precise, the equilibrium is not first-order stable, so that the perturbation solution method is not feasible. Global stability could still hold in principle.

Figure 2.11: Parameter sensitivity.

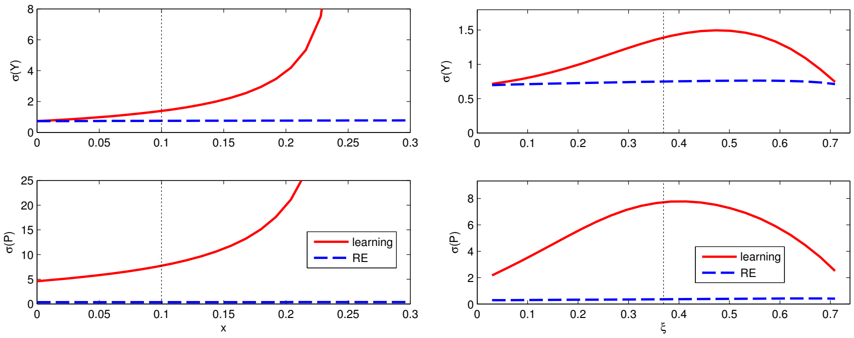

(a) Market value dependency x. (b) Recovery rate ξ.

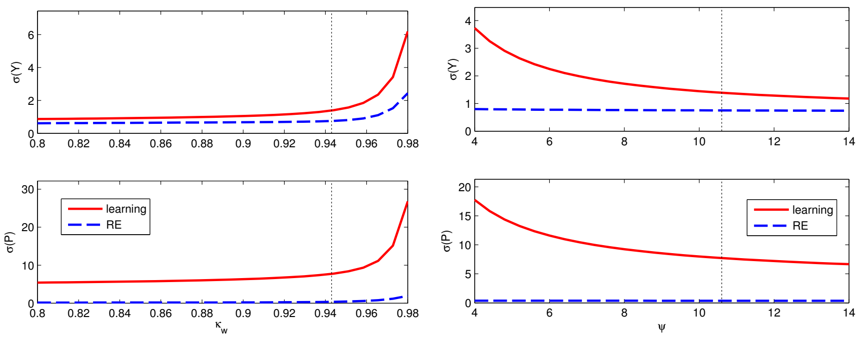

(c) Wage rigidity κ_w. (d) Investment adjustment cost ψ.

Theoretical standard deviations, unfiltered. Dashed black line indicates parameter value in the estimated learning model.

the rational expectations solution barely depends on the parameter x. This might be one reason why the distinction between market and liquidation value has not featured prominently in the existing literature on firm credit frictions.

Panel (b) shows the dependency on the average tightness of credit frictions ξ. Amplification is hump-shaped with respect to this parameter. At $\xi = 0$, no collateral is pledgeable and firms cannot borrow at all. In this case, fluctuations in stock prices do not matter and learning does not introduce any amplification. On the other hand, as pledgeability increases to its maximum value (beyond which a steady state with permanently binding borrowing constraint does not exist), amplification also disappears. This mirrors the analysis of the simplified model: when borrowing constraints relax, general equilibrium effects in the form of wage and interest rate changes offset amplification.

Panels (c) and (d) document the importance of wage rigidity and investment adjustment costs, respectively. Wage rigidity helps to improve the fit of the model, but at the same time it also helps amplification. When credit constraints relax, a larger degree of rigidity mitigates the negative general equilibrium effect on firm profits, raising internal net worth and stock prices. In contrast, adjustment costs dampen amplification since they make investment more costly. The sensitivity to these parameters is much greater under learning than under rational expectations.

2.7 IMPLICATIONS FOR MONETARY POLICY

In the model with learning, changes in subjective expectations in financial markets lead to large and inefficient asset price and business cycle fluctuations. A natural question, therefore, is whether policy should intervene in order to stabilise asset prices. The question is not so much whether a monetary policy reaction to asset prices is desirable at all in theory, but whether the benefits are predicted to be sizeable enough to warrant the inclusion of such a volatile indicator into a monetary policy framework in practice. This question has been discussed for a long time by academics and policymakers alike, but generally answered in the negative. However, the discussion has generally centred on models in which asset price fluctuations are either efficient, or the inefficient ("bubble") component cannot be reduced by raising interest rates. However, when expectations in financial markets are not rational, the potential benefits of reacting to asset prices are much larger. There is a possibility that the central bank can mitigate excessive fluctuations in subjective expectations, thereby substantially stabilising the business cycle.

The model does not permit to solve analytically for optimal monetary policy, but I can numerically evaluate the effect of a class of interest rate rules, augmented with a reaction to asset prices. Consider extending the interest rate rule (2.41) as follows:

$$i_t = \rho_i i_{t-1} + (1-\rho_i)\begin{pmatrix} 1/\beta + \pi^* + \phi_\pi\left(\pi_t - \pi^*\right) \\ +\phi_{\Delta Y}\left(\log Y_t - \log GY_{t-1}\right) \\ +\phi_{\Delta P}\left(\log P_t - \log GP_{t-1}\right) \end{pmatrix} \tag{2.63}$$

In addition to raising interest rates when inflation is above its target level, the central bank can raise interest rates by $\phi_{\Delta Y}$ percentage points when real GDP growth is above long-run TFP growth, and by $\phi_{\Delta P}$ percentage points when stock market growth is above long-run TFP growth. It also rules out monetary policy shocks. I explicitly do not include the levels of output or asset prices or output gap measures. From a theoretical perspective, this would imply that the central bank has more knowledge than the private sector under learning, since the long-run level of asset prices and output is believed to be essentially a random walk. From a practical perspective, imposing a level target for asset prices is an even more audacious measure than a target for price growth.

I compute the parameters for which such a rule maximises conditional welfare (Schmitt-Grohe and Uribe, 2004) (evaluated as a second-

Table 2.6: Optimal policy rules.

	learning			RE		
	(4) baseline	(5) w/o ΔP	(6) w/ ΔP	(1) baseline	(2) w/o ΔP	(3) w/ ΔP
ϕ_π	1.5	1.27	1.55	1.50	5.03	4.81
$\phi_{\Delta Y}$		.13	.19		-.14	.21
$\phi_{\Delta P}$			.40			.18
ρ_i	.5	.12	.04	.50	.99	.99
$\sigma(Y)$	3.31	3.32	2.55	2.39	2.13	2.13
$\sigma(P)$	10.18	7.20	4.85	.215	.401	.401
$\sigma(\pi)$	.401	.292	.249	.393	.177	.178
$\sigma(i)$	.434	.462	.225	.424	.050	.052
%c. loss	**.741**	**.309**	**.028**	**.134**	**.0246**	**.0246**

Standard deviations of output, stock prices, inflation, and interest rates (unfiltered) in percent. Welfare loss based on second order approximation to objective expected conditional welfare.

order approximation around the non-stochastic steady-state). Under learning, I use objective, i.e. statistical expected welfare, rather than expected welfare under subjective beliefs $\mathcal{P}$. This criterion is paternalistic in the sense that it does not coincide with the policy that agents in the model would prefer the central bank to take. In this sense, I assume that the central bank is able to commit to a rule that goes against what markets think it should do.

Table 2.6 summarises the key findings. Column (1) reports the baseline model under learning. The bottom row shows the welfare criterion in terms of equivalent steady-state consumption loss. This consumption loss is .74% under learning and an order of magnitude larger than the usual Lucas cost of business cycles. This is partly due to inflation fluctuations with Calvo price setting, but partly due to the fact that under learning, agents make choices that are optimal under subjective beliefs but highly suboptimal when evaluated under the objective probability measure. Column (2) calculates the welfare-maximising rule that does not include a reaction to asset prices ($\phi_{\Delta P} = 0$). This makes it possible to cut the welfare loss by more than half, to .31%; however, output and stock price volatility remain high. A reaction to stock price growth, however, is able to deliver substantial additional stabilisation. The coefficient $\phi_{\Delta P}$ is large and positive: A quarter-over-quarter increase in stock prices by 1% is met with an interest rate rise of almost 0.4%. This does not mean, however, that the central bank interest rate is itself very volatile. In fact, the standard deviations of output, stock prices, and interest rates are considerably reduced. By raising interest rates in a stock price boom, the central bank lowers aggregate demand and raises the cost of borrowing for firms, which both act to reduce corporate profits. This curbs investor optimism in the stock market, thus dampening the feedback loop between beliefs, prices, and profits.

This result is specific to the learning equilibrium and is not obtained under rational expectations, as shown in Columns (4) to (6). When learning is shut down, the welfare cost of business cycles for the baseline parameterisation is already much lower to start with. Optimising over the reaction to inflation and output growth alone, the central bank can already effectively stabilise the economy and reduce the welfare loss considerably. Column (6) shows that when $\phi_{\Delta P}$ is unconstrained, then the optimal rule is calculated to include a positive reaction to asset price growth as well. However, the added benefit of doing so is negligible, with the welfare cost virtually unchanged. Hence, under rational expectations, a monetary policy reaction to asset prices carries no benefits.

2.8 Conclusion

This chapter has analysed the implications of a learning-based asset pricing theory in a business cycle model with financial frictions. When firms borrow against the market value of their assets, learning in the stock market interacts with credit frictions to form a two-sided feedback loop between beliefs, stock prices and firm profits that amplifies the asset price dynamics.

I have embedded the mechanism in a dynamic stochastic general equilibrium model with nominal rigidities. Unlike most of the literature on adaptive learning, beliefs retained a high degree of rationality and internal consistency. In the baseline calibration of the model, introducing learning was shown to considerably improve the model's asset price properties, notably volatility and return predictability, while still matching standard business cycle statistics. At the same time, it leads to a large amount of propagation and amplification of both supply and demand shocks, endogenising up to one third of the volatility of the business cycle.

A natural criticism of a theory based on beliefs other than rational expectations is that it has many degrees of freedom to adjust for almost any fact of choice. Here, expectations were constructed such that learning only adds one additional parameter, the learning gain g, which is calibrated to the relative volatility of stock prices over dividends. Moreover, the predictable biases in forecasts made by agents in the model correspond surprisingly well to the patterns found in actual data from surveys.

The model was also used to study normative implications of learning. In particular, I have revisited the question of whether monetary policy should react systematically to asset prices. I found that a strong reaction to stock price growth is desirable from a welfare perspective when investors in financial markets are learning. In contrast, under rational expectations, such a reaction does not improve welfare, in line with previous findings in the literature. This illustrates the importance of asset pricing for macroeconomics. Assumptions about the source of volatility in financial markets has profound consequences for our policy recommendations.

While this chapter mainly argues that improving asset price dynamics in a macroeconomic model can lead to novel insights about business cycles, future work should explore the other direction: Can a model of learning in a production economy teach us something new about asset pricing? I suspect that the endogenous feedback between beliefs and real activity might provide for insights into the joint predictability of dividend growth and returns. It could also be worthwhile to explore extensions to the learning mechanism (such as a different perceived law of motion for prices, or learning about returns) that might further improve the asset price characteristics, a greater persistence of the P/D ratio and possibly even account for the equity premium.

But the analysis of this model need not be restricted to the role of stock prices affecting investment. Prices of other assets, for example housing, are also known to exhibit patterns that are difficult to reconcile with the simple asset pricing models used commonly in the macro literature. At the same time, they clearly play a crucial role in affecting business cycle fluctuations through credit constraints. A learning-based approach is likely to be fruitful in this context.

3 OPTIMAL UNEMPLOYMENT INSURANCE AND INTERNATIONAL RISK SHARING

3.1 INTRODUCTION

Europe has recently seen business cycle movements differ greatly across countries. This development, together with the resulting strains on public budgets, has renewed calls to introduce some form of public cross-country risk sharing, sometimes under the name of a "fiscal union". Indeed, a widely held view is that a common currency exacerbates the need for international risk sharing mechanisms, and that fiscal transfers become desirable when the private sector lacks such mechanisms (Mundell, 1961). Capital and labour markets are not nearly as integrated across member countries in the European Monetary Union as they are across states in the United States, so this is a relevant concern for the Eurozone.

At the same time, high unemployment levels in many developed countries have led to renewed interest in the design of unemployment insurance. In the Eurozone in particular, policy makers have argued that the unemployment insurance system is a good and politically viable channel to share risk across countries.[1]

This chapter asks the question: If a group of countries were to introduce a common unemployment insurance system, what should it look like? We answer this question using a two-country business cycle model with search frictions in labour markets. Financial markets are incomplete and labour is immobile across countries, so that country-specific risk and idiosyncratic unemployment risk can only be partially insured privately. The government in each country maintains a mandatory unemployment insurance system providing benefits to unemployed workers and taxing employed workers. On top of this, a supranational unemployment insurance agency is introduced which is able to administer an additional component of the unemployment insurance system. This component can vary across countries as a function of country-specific shocks.

1 A harmonised unemployment insurance system within the Eurozone as a tool for international risk sharing has been suggested by the President of the European Council (van Rompuy, 2012), the International Monetary Fund (Blanchard et al., 2014), the German Institute for Economic Research (Bernoth and Engler, 2013) and the French Advisory Council (Artus et al., 2013). Brenke (2013) also discusses some of the drawbacks.

We derive two theoretical insights from our analysis. Our first result is that a supranational unemployment insurance system can in principle be used to insure against country-level risk without affecting unemployment levels. The intuition for this result is as follows. Unemployment insurance affects unemployment levels by changing the relative value of employment over unemployment, which determines incentives to search and wage bargaining outcomes. When a country is to receive a fiscal transfer, this relative value can be kept constant by simultaneously increasing the level of benefits and lowering the rate of contributions to the unemployment insurance system. The opposite can be done in the country which is to send the transfer.

Next, we are interested in how cross-country transfers interact with optimal labour market policies. Our second result is that the presence of an international risk sharing motive introduces a countercyclical element to the optimal unemployment insurance policy. The intuition is as follows. The classic unemployment insurance tradeoff for a social planner is between efficiency of employment and insurance. Too much insurance reduces search effort and/or job creation, while too little insurance harms risk-averse workers who cannot insure privately against unemployment risk. When international risk sharing is present, the planner is shielding local consumption from fluctuations in local output. When output falls, its share in local consumption falls as well. After a negative productivity shock in one country, the planner can then afford to provide more generous insurance and shift employment towards countries where it is more productive. Therefore, insurance becomes more countercyclical than without cross-country transfers.

We calibrate our model to the Euro area and compute the Ramsey-optimal policy. We then compare this with a policy of constant replacement rates and no international transfers (as they are currently in place in most countries). Our baseline simulation suggests that the optimal unemployment insurance policy is countercyclical, even when nationally optimal policies would prescribe procyclicality. The optimal policy is very far from the constant replacement rate benchmark and involves sizeable changes in replacement rates and transfers over the business cycle. We also compare the optimal policy to a recent proposal by Artus et al. (2013) to implement a EU-wide unemployment insurance. We find that in contrast to their findings, the proposal only has very small effects in our model. The main reason is that their calculations do not take into account the surpluses and deficits of their scheme which ultimately have to be financed by national governments. When we take these fiscal effects into account in our general equilibrium model, the effectiveness of the proposal is greatly reduced.

Our results have several limitations. First, the only relevant sources of employment fluctuations in our model are productivity shocks.

The first-best outcome in this case, well known from the international RBC literature, is to reduce employment in a country experiencing a negative shock. In competitive equilibrium, this can be realised by making unemployment insurance more generous. But it is not clear that the same would be true for other types of shocks. In this sense, our analysis is a first pass at the policy design problem. Second, we abstract from sovereign debt and impose balanced budgets on national governments. When cyclical fluctuations are sufficiently short-lived, governments can effectively implement risk-sharing by financing countercyclical fiscal policy with debt, and it is not clear that the benefits of a supranational risk-sharing scheme are large. However, recent events in the Eurozone have made it clear that there are limits on the debt capacity of governments, especially at times when risk sharing might be needed most. Our assumption of balanced budgets can be interpreted as an extreme form of a sovereign debt constraint. Third, we abstract from the political moral hazard induced by risk-sharing (Persson and Tabellini, 1996). It is plausible that a fiscal transfer mechanism reduces incentives for national governments to carry out structural reforms and this is probably the main political reason for its opposition in the Eurozone. But ultimately, such concerns need to be weighed against the economic benefits of the mechanism. In this chapter, we are focusing on the latter.

The remainder of the chapter is organised as follows. We briefly review the related literature in Section 3.2. In Section 3.3, we lay out a simplified version of our model with only two periods. This allows us to show our theoretical insights analytically and provide intuition. In Section 3.4, we lay out the full dynamic model that we use for quantitative analysis and calibrate it to the Euro area. Section 3.5 contains the numerical results from our calibrated model. Section 3.6 concludes.

3.2 RELATED LITERATURE

Our analysis relates to the literature on international risk sharing and fiscal union on one hand, and the literature on the design of optimal unemployment insurance on the other hand.

It is well known that the search externality in frictional labour markets can be corrected using unemployment insurance. Because of costly search, employment – and the corresponding fluctuations – may be too low or too high, depending mainly on the relation of the workers' bargaining power to the matching elasticity. In the steady state, this can be resolved by changing the outside option of workers through unemployment benefits (Hosios, 1990). When workers are risk-averse, the correction of the search externality needs to be weighed against the provision of insurance (Baily, 1978). Fredriksson and Holmlund (2006) provide surveys of the literature on optimal unemployment insurance in static and steady-state situations. More recently, interest has emerged in unemployment insurance policies that depend on the state of the business cycle. Here, a central point of debate is whether benefits should become more generous in a recession in order to increase insurance (countercyclical policy), or less generous in order to mitigate the fall in employment (procyclical policy). Earlier contributions such as Kiley (2003) and Sanchez (2008) suggest that there is room for countercyclical unemployment benefit policy. A more recent contribution is Landais et al. (2015). They analyse a model with sticky wages and job rationing and find that a countercyclical policy is optimal as the effects of insurance on equilibrium unemployment is smaller in recessions. On the other side, Mitman and Rabinovich (2015) numerically compute optimal dynamic policies and show that the cyclical stance of the unemployment insurance is procyclical, in a setting when workers' outside option leads to inefficiently high wages. Moyen and Stähler (2014) analyse the optimal cyclicality holding its average level fixed. They show that there are situations in which unemployment insurance should be countercyclical even when wages are directly affected and the bargaining power of workers is too high relative to the Hosios condition.[2] Jung and Kuester (2015) analyse first-best policy with sufficiently many fiscal instruments. They find that, in recessions, benefits should rise if, at the same time, hiring subsidies and layoff taxes also rise. The latter two instruments increase the incentives to hire and decrease those to fire workers, which may compensate partly for increased wage costs. However, if the other two instruments are not available, they also find procyclical benefits to be optimal.

The literature exclusively analyses closed economies. This chapter analyses optimal policy in a context in which unemployment insurance operates across multiple countries and faces the additional objective of sharing cross-country risk.

Turning to the literature on fiscal unions, Leduc et al. (2009) have shown that, when asset markets are incomplete, country-specific productivity disturbances can have large uninsurable effects on wealth and consumption paths. In a prominent recent paper, Farhi and Werning (2012) find that such uninsurable effects may be especially large in a currency union with nominal rigidities. They suggest forming a transfer union to insure against this risk. Many economists follow their view that, in federal unions, a (fiscal) transfer mechanism

2 Moyen and Stähler (2014) compare optimal benefit duration policy in Europe and the US. In their European calibration, the bargaining power of workers is larger than the matching efficiency, implying the optimal benefit to be negative in light of the Hosios condition. However, it is restricted to be positive. Additionally, rule-of-thumb households make average marginal utility of consumption fluctuate relatively much. It can be shown that steady-state benefits above optimum and relatively volatile marginal utility of consumption makes optimal benefit policy countercyclical even when bargaining power of workers is high already.

to at least compensate for the uninsurable effects due to nominal rigidities may be desirable. However, there is still some debate on how to ideally establish such a transfer mechanism or a fiscal union (see Bargain et al., 2013 and Bordo et al., 2011 for a discussion). In this chapter, we show that a transfer mechanism is even desirable in a model without nominal rigidities. Including them should only strengthen our results.

3.3 SIMPLIFIED MODEL

The intuition for our results can best be seen in a simple two-period model which allows us to analytically prove our results and provide a graphical representation. The quantitative analysis is carried out in the next section.

3.3.1 *Model setup*

There are two countries, Home and Foreign. The Home country is inhabited by a mass $\omega \in (0, 1)$ of workers, while the Foreign country is inhabited by a mass $1 - \omega$ of workers. In each country, firms transform labour into consumption goods. Firms are owned by risk-neutral entrepreneurs. While consumption goods can be traded across countries in competitive markets, labour is immobile across countries and labour markets are subject to search frictions.

In the first period, all workers start out as unemployed and no production takes place. Agents can, however, trade assets with each other. Asset markets are incomplete, and we will spell out the precise market structure later on. In any event, the utility function of a worker in Period 1 is as follows:

$$U = \mathbb{E}\left[nu\left(c_e\right) + \left(1 - n\right) u\left(c_u\right)\right] \tag{3.1}$$

where c_e is his consumption level if he turns out to be employed in Period 2, and c_u his consumption level when he turns out to be unemployed. n is the employment level in Period 2. We assume logarithmic utility, $u\left(c\right) = \log\left(c\right)$.

In the second period, firms post vacancies, workers are matched with firms and production takes place. In the Home country, the initial unemployment rate is $u = 1$ and the number of vacancies is v.[3] The number of matches follows a Cobb-Douglas production function

$$m\left(u, v\right) = \kappa_m u^{\mu} v^{1-\mu}. \tag{3.2}$$

Since the past stock of employment is zero, employment at the end of the period is

$$n = m\left(1, \theta\right) \tag{3.3}$$

3 Throughout the chapter, quantities will be expressed in per capita terms unless otherwise indicated.

where $\theta = v/u$ is labour market tightness.

A firm that posts a vacancy incurs a cost κ_v. The probability that the vacancy is filled is $q(\theta) = \kappa_m \theta^{-\mu}$. In that case, the match produces output a and the worker gets paid a wage w. This wage is determined using Nash bargaining, where the bargaining power of workers is denoted ξ (the bargaining solution is described further below). A zero-profit condition for vacancy creation prescribes

$$\kappa_v = q(\theta)(a - w) \tag{3.4}$$

We denote by y aggregate output in the Home country net of vacancy costs:

$$y = an - \kappa_v v \tag{3.5}$$

The productivity a is a random variable which is only revealed in the second period.

Employed workers receive wages w which are taxed at the rate τ, while unemployed workers receive unemployment benefits b. Each worker might also receive income W_i from assets traded in the first period, where $i = e, u$ denotes his status as employed or unemployed. The individual and aggregate consumption levels are:

$$c_e = (1 - \tau) w + W_e \tag{3.6}$$

$$c_u = b + W_u \tag{3.7}$$

$$c = nc_e + (1 - n) c_u \tag{3.8}$$

The Foreign country has a similar structure to the Home country, but with possibly different parameters. We denote Foreign variables with an asterisk, e.g. b^* for foreign unemployment benefits. Home and foreign productivity (a, a^*) are the only sources of aggregate uncertainty.

Payroll taxes τ and benefits b are administered by an unemployment insurance agency. We assume that the two countries are part of an insurance union, such that the agency operates across both countries. It has to run a balanced budget with the constraint:

$$\omega [(1 - n) b - n\tau w] + (1 - \omega) [(1 - n^*) b^* - n^* \tau^* w^*] = 0. \tag{3.9}$$

In order to close the model, we have to specify the assets that agents can use in the first period to insure themselves against risk, and the unemployment insurance policies.

3.3.2 *Social planner solution*

Before looking at the competitive equilibrium, we first look at a benchmark social planner solution. A utilitarian social planner maximises the sum of worker utilities subject only to the resource constraint and the search friction by solving the following problem:

$$\max_{\begin{pmatrix} n,\theta,c^e,c^u \\ n^*,\theta^*,c^{e*},c^{u*}\end{pmatrix}} \begin{array}{ll} \tilde{\omega} & \mathbb{E}\left[nu\left(c_e\right)+\left(1-n\right)u\left(c_u\right)\right] \\ +\left(1-\tilde{\omega}\right) & \mathbb{E}\left[n^*U\left(c_u^*\right)+\left(1-n^*\right)U\left(c_u^*\right)\right] \end{array}$$

$$\begin{aligned} \text{s.t. } n &= \kappa_v\theta^{1-\mu} && (3.10)\\ n^* &= \kappa_v^*\left(\theta^*\right)^{1-\mu^*} && (3.11)\\ \omega\left(nc_e+\left(1-n\right)c_u\right) & \\ +\left(1-\omega\right)\left(n^*c_e^*+\left(1-n^*\right)c_u^*\right) &= \omega\left(an-\kappa_v\theta\right) \\ &+ \left(1-\omega\right)\left(a^*n^*-\kappa_v^*\theta^*\right) && (3.12)\end{aligned}$$

Here, $\tilde{\omega}$ is the relative weight the planner puts on workers in the Home country, which might be more or less than the size of its population ω. Within a country, all workers are ex-ante homogenous and so weighting of individual workers is inconsequential.[4] The first order conditions of the planner problem are standard:

$$\begin{aligned} \kappa_v &= \kappa_m\theta^{-\mu}\left(1-\mu\right)a && (3.13)\\ \kappa_v^* &= \kappa_m^*\left(\theta^*\right)^{-\mu^*}\left(1-\mu^*\right)a^* && (3.14)\\ c_u &= c_e && (3.15)\\ c_u^* &= c_e^* && (3.16)\\ \frac{\omega}{\tilde{\omega}}c_e &= \frac{1-\omega}{1-\tilde{\omega}}c_e^* && (3.17)\end{aligned}$$

The first two conditions are the Hosios conditions in each country, which determine the number of vacancies that maximise aggregate output net of vacancy costs. The remaining conditions prescribe full risk sharing within and across countries. The consumption levels of employed and unemployed workers within each country should be identical, and each country should consume a constant fraction of union output.

3.3.3 *Optimal policy with private insurance*

We now come back to the competitive equilibrium. Even when markets are complete, the competitive equilibrium generally doesn't implement the social planner solution because of search externalities. Throughout this chapter, we assume some form of market incompleteness, since our focus is on how unemployment insurance can be used to overcome insufficient international risk sharing. In this section,

4 The entrepreneurs owning firms make zero profits in the decentralised equilibrium in all states of the world. Following the logic of Landais et al. (2015), we constrain the planner to keep entrepreneurial consumption at zero as well.

we allow workers to only insure domestically against idiosyncratic unemployment risk. In Period 1, each worker $i \in [0, \omega]$ at Home can issue a claim on his future income. These claims can be traded in a competitive market within the country but not across countries. Since all workers are ex-ante identical, it is optimal for a Home worker to fully diversify his risk by selling his entire future income in exchange for a diversified portfolio of the income of all other Home workers' income. In this case, the consumption levels in Period 2 are

$$c_e = c_u = c \tag{3.18}$$

This allows us to solve for the Nash-bargained wage. The worker surplus from a match is

$$\mathcal{W} - \mathcal{U} = u'(c)\frac{\partial c}{\partial n} \tag{3.19}$$

and the firm surplus is

$$\mathcal{J} = a - w. \tag{3.20}$$

When workers have bargaining power ξ, the bargained wage is simply:

$$w(a, \rho) = \frac{\xi a}{\xi + (1-\xi)(1-\rho)} \tag{3.21}$$

where ρ is the net replacement rate, defined as

$$\rho = \frac{b}{(1-\tau)w}. \tag{3.22}$$

A higher replacement rate improves workers' outside option and drives up wages. It thereby lowers the incentives for job creation and reduces employment.

We want to know what the optimal unemployment insurance scheme looks like in this situation. With privately insured unemployment risk, the insurance agency has to mitigate three inefficiencies: search externalities in the Home and Foreign country and lack of international risk sharing. It also has three policy instruments: the Home and Foreign replacement rates and a cross-country transfer. This already suggests that there exists a policy that eliminates all inefficiencies.

We first note that the budget constraint of the unemployment insurance agency can be rewritten as

$$\begin{aligned} 0 = \quad & \omega\left[(1-n)b - n\tau w\right] + (1-\omega)\left[(1-n^*)b^* - n^*\tau^* w^*\right] \\ = \quad & \omega(c-y) + (1-\omega)(c^* - y^*) \end{aligned} \tag{3.23}$$

in which the replacement rates do not appear. We can therefore choose replacement rates ρ, ρ^* and a cross-country transfer $\omega(c-y)$ as a policy, and back out the necessary benefits b, b^*, τ, τ^* from the budget constraint and replacement rate definition.

Unemployment rates only depend on policy through the replacement rate. Therefore, *a transfer of resources from one country to another can be implemented through the unemployment insurance system without affecting unemployment levels*. A positive transfer from Foreign to Home would be implemented by increasing benefits b to unemployed workers, and at the same time lowering payroll taxes τ on employed workers. This way, all workers get to consume more, but the net replacement rate ρ stays constant and the relative bargaining position of workers is unchanged.[5]

The replacement rates satisfying the Hosios condition are

$$\rho = \frac{\mu - \xi}{\mu (1 - \xi)}, \rho^* = \frac{\mu^* - \xi^*}{\mu^* (1 - \xi^*)} \tag{3.24}$$

and the optimal consumption with a social planner weight $\tilde{\omega}$ on the Home country is

$$c = \frac{\tilde{\omega}}{\omega} (\omega y + (1 - \omega) y^*) . \tag{3.25}$$

The planner weight $\tilde{\omega}$ can be chosen freely. Here, we determine it by the condition that transfers are zero in expectation, so that in Period 1, neither country expects to be subsidising the other country on average through the unemployment insurance system. Imposing $\mathbb{E}[c - y] = 0$ leads to a planner weight that is simply the expected share of Home output in union output

$$\tilde{\omega} = \frac{\mathbb{E}[\omega y]}{\mathbb{E}[\omega y + (1 - \omega) y^*]} \tag{3.26}$$

and a transfer policy

$$c - y = \frac{\mathbb{E}[y] \, \mathbb{E}[y^*]}{\mathbb{E}\left[\frac{\omega}{1-\omega} y + y^*\right]} \left(\frac{y}{\mathbb{E}[y^*]} - \frac{y}{\mathbb{E}[y]} \right) . \tag{3.27}$$

The Home country receives a transfer when its output is below average, but has to pay a transfer when output in the Foreign country is below average. This policy perfectly replicates the social planner solution.

3.3.4 *Optimal policy without private insurance*

The previous case has illustrated how the unemployment insurance system can implement cross-country transfers orthogonally to unemployment levels. However, we have so far abstracted from the most important objective of unemployment insurance, namely to insure

5 This is a general result: An unemployment insurance scheme which can vary benefits and contributions in both countries has four policy instruments and one budget constraint, which means that it is possible to achieve three objectives, in particular leaving employment levels unchanged while implementing a cross-country fiscal transfer.

against unemployment. In the presence of privately uninsurable unemployment risk, the optimal policy becomes genuinely second-best and tradeoffs emerge between all three policy objectives: maximising net output, providing insurance between employed and unemployed, and providing insurance across countries.

We now eliminate all asset trade in Period 1, which means that workers cannot insure any risk. The consumption levels in Period 2 are simply:

$$c_e = (1-\tau)\, w \tag{3.28}$$

$$c_u = b = \rho c_u \tag{3.29}$$

We solve again for the Nash-bargained wage, which now takes into account the curvature in the worker's utility function. The worker surplus from a match is

$$\mathcal{W} - \mathcal{U} = u\,(c_e) - u\,(c_u) \tag{3.30}$$

and the firm surplus is unchanged. When workers have bargaining power ξ, the bargained wage is now:

$$w\,(a,\rho) = \frac{\xi a}{\xi - (1-\xi)\log\rho}. \tag{3.31}$$

In this situation, the social planner allocation is no longer feasible. Providing full insurance against idiosyncratic unemployment risk clearly calls for $\rho = 1$, but in this case the worker gets to capture the whole surplus ($w = a$) and job creation completely collapses. Therefore, we have to solve for the Ramsey-optimal policy here.

The Ramsey planner solves:

$$\max_{\substack{n,\theta,c^e,\rho,\\ n^*,\theta^*,c^{e*},\rho^*}} \quad \begin{array}{rl} \tilde{\omega} & \mathbb{E}\left[nu\,(c_e) + (1-n)\,u\,(\rho c_e)\right] \\ +\,(1-\tilde{\omega}) & \mathbb{E}\left[n^* u\,(c_e^*) + (1-n^*)\,u\,(\rho^* c_e^*)\right] \end{array}$$

$$\begin{aligned} \text{s.t. } n &= \kappa_m \theta^{1-\mu} && (3.32)\\ n^* &= \kappa_m^*\,(\theta^*)^{1-\mu^*} && (3.33)\\ \kappa_v &= \kappa_m \theta^{-\mu}\,(a - w\,(a,\rho)) && (3.34)\\ \kappa_v^* &= \kappa_m^* \theta^{*-\mu^*}\,(a^* - w^*\,(a^*,\rho^*)) && (3.35)\\ \omega\,(n + (1-n)\,\rho)\,c_e + \\ (1-\omega)\,(n^* + (1-n^*)\,\rho^*)\,c_e^* &= \omega w\,(a,\rho)\,n \\ &+ \ (1-\omega)\,w^*\,(a^*,\rho^*)\,n^* && (3.36) \end{aligned}$$

Here, we have substituted out many of the equilibrium conditions of the competitive equilibrium. In particular, choosing an unemployment insurance policy (b, b^*, τ, τ^*) subject to the insurance agency's budget constraint is equivalent to choosing Home and Foreign replacement rates and consumption levels $(\rho, \rho^*, c_e, c_e^*)$ subject to the aggregate resource constraint. As before, we choose the social planner $\tilde{\omega}$ such that any transfers made across countries net out in expectation: $\mathbb{E}\left[c-y\right]=0$.

As we have written it, the problem has eight choice variables and five constraints, leaving three degrees of freedom. These correspond of course to the three policy instruments ρ, ρ^* and the cross-country transfer $c-y$. The first order condition determining the optimal transfer is as follows:

$$c-y=\frac{\mathbb{E}\left[y\right]\mathbb{E}\left[y^*\right]}{\mathbb{E}\left[\frac{\omega}{1-\omega}y+y^*\right]}\left(\frac{y^*}{\mathbb{E}\left[y^*\right]}-\frac{y}{\mathbb{E}\left[y\right]}\right). \tag{3.37}$$

This is the exact same condition as in the previous case: Each country at optimum consumes a constant share of union output. The Home country receives a transfer when its output is below average, but has to pay a transfer when output in the Foreign country is below average.[6] However, this now only holds for *average* consumption in a country. The consumption of each worker need not be proportional to union output.

The central equation in this section is the first order condition with respect to the replacement rate. For the Home country, it reads as follows:

$$\underbrace{\frac{(1-n)(1-\rho)}{n+(1-n)\rho}-\epsilon_\rho^n\left(\log\rho+\frac{1-\rho}{n+(1-n)\rho}\right)}_{=:I(\rho)}=\underbrace{-\epsilon_\rho^y\frac{1}{n}\frac{y}{c}}_{=:H(\rho)} \tag{3.38}$$

where $\epsilon_\rho^n=\frac{dn}{d\rho}\frac{\rho}{n}$ is the elasticity of Home employment with respect to the Home replacement rate, and $\epsilon_\rho^y=\frac{dy}{d\rho}\frac{\rho}{y}$ is the elasticity of net Home output with respect to the Home replacement rate. A symmetric condition is obtained for the Foreign country.

This condition has an intuitive interpretation. The left-hand side, which we call $I(\rho)$, is the marginal benefit of insurance when raising the replacement rate, at a fixed quantity of output available to the country. By raising ρ, the unemployed's marginal utility increases relative to average marginal utility. This is the first term on the left-hand side of Equation (3.38). At the same time, a higher ρ reduces employment (through higher wages and lower job creation) which shifts the composition of the workforce towards the unemployed. This means that one marginal worker suffers a utility loss, which is the "$\log\rho$" term in the left-hand side of Equation (3.38). It also implies a composition effect on the insurance budget, captured by the remaining term on the left-hand side. The right-hand side, which we call

6 This result is due to our assumption of logarithmic utility.

Figure 3.1: Optimal replacement rate for constant a and y/c.

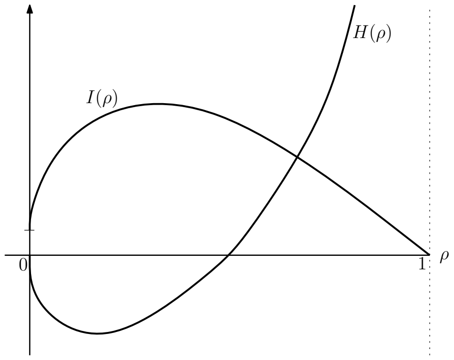

$H(\rho)$, is the marginal cost of raising the replacement rate in terms of net output lost (output minus vacancy costs).

The determination of the optimal replacement rate is graphically depicted in Figure 3.1, which plots the functions $H(\rho)$ and $I(\rho)$.[7] We can see that the insurance term $I(\rho)$ is positive and only equals zero at $\rho = 1$. Intuitively, holding output constant it is always desirable to increase the replacement rate until full insurance is achieved. The efficiency term $H(\rho)$ is first negative and then turns positive, approaching plus infinity at $\rho \to 1$. Intuitively, when ρ is too high, there is too little job creation and the amount of resources available for consumption can be increased by lowering replacement rates, thereby lowering bargained wages and increasing job creation. In this case, $\epsilon_\rho^y < 0$ and therefore $H(\rho)$ is also positive. As $\rho \to 1$, output collapses to zero and the marginal utility from lowering the replacement rate becomes infinite. Conversely, when ρ is too low, there is too much vacancy posting and the amount of resources available for consumption can be increased by raising replacement rates. In this case, $H(\rho)$ is negative.

The optimal replacement rate lies at the intersection between the two curves. We can already see that under the optimal policy, employment is always lower than in the social planner solution. Since the benefit of insurance is positive, the optimal ρ is always higher than what the Hosios condition $H(\rho) = 0$ would call for.

What happens to the optimal replacement rate when shocks to a or a^* hit the economy? We first keep the ratio y/c constant (one can imagine a closed-economy situation in which $y/c = 1$) and look at the effect of a reduction in productivity a. The effects are depicted in in Panel (a) of Figure 3.2.

A reduction in a increases the insurance term $I(\rho)$ and scales up the efficiency term $H(\rho)$. The intuition is as follows. In this model, a reduction in a has a negative effect on employment because vacancy creation costs κ_v are fixed. Holding total resources constant, the de-

7 Proposition A.2 in the appendix proves that the shape of the I and H curves are indeed as depicted.

Figure 3.2: Change of the optimal replacement rate in a recession.

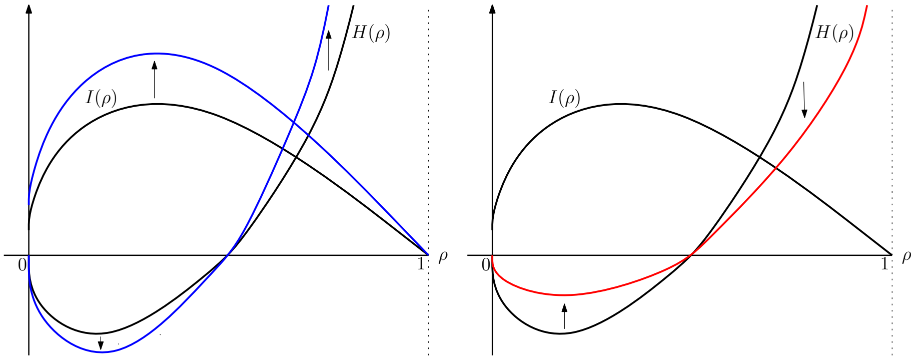

(a) Reduction in a. (b) Reduction in y/c.

crease in employment translates into an increase in unemployment risk for workers, raising the social benefit to insure. Therefore, $I(\rho)$ shifts up for any value of ρ. At the same time, lower productivity directly reduces net output for any level of employment. Therefore, the average marginal utility of increasing output towards its efficient level increases and $H(\rho)$ is scaled up for any value of ρ. These two forces work in opposite directions on the replacement rate, so that the overall effect is ambiguous.

So far, we have kept the ratio y/c constant, but at optimum it is jointly determined with the replacement rate. The risk-sharing condition (3.37) prescribes that average consumption c of the Home country is proportional to union output. If Home produces more, then consumption rises less than one-for-one: y/c is increasing in y and decreasing in y^*. It is this risk-sharing aspect which is novel to the literature on optimal unemployment insurance.

The presence of international risk-sharing makes the replacement rate more countercyclical. This can be seen easily from the optimality condition (3.38). When a falls, Home's output will relatively low compared to union output, and y/c will drop. Panel (b) of Figure 3.2 shows that the optimal replacement rate rises as a response, introducing a countercyclical element to the optimal policy. When y/c falls, the efficiency term $H(\rho)$ gets compressed towards zero, which can be seen directly from Equation (3.38). Intuitively, Home's output is now relatively less important compared to union output. Therefore, it becomes less important to ensure that this output is at its efficient level. The tradeoff between efficiency and insurance shifts towards the latter and the replacement rate becomes more generous.[8]

We can show (see Proposition A.5 in the appendix) that *in the limit as $\omega \to 0$, the Home replacement rate when y/c is varied optimally is decreasing in Home productivity a (countercyclical)*. The smaller ω, the better Home country risk can be hedged and as the country's size

8 Proposition A.4 in the appendix provides a formal proof.

approaches zero relative to the entire risk-sharing union, average consumption can be completely shielded from output fluctuations. In this case, a fall in a unambiguously raises the replacement rate. Nevertheless, the Ramsey planner's tradeoff between efficiency and insurance does not disappear – the optimal replacement rate always remains below one.

Also, *the Home replacement rate is increasing in Foreign productivity* a^*, as this variable affects Equation (3.38) only through a lower ratio y/c. When the Foreign country experiences a drop in productivity, maximising Home output now matters relatively more for the Ramsey planner, and the tradeoff between efficiency and insurance shifts towards the former. The Home replacement rate therefore becomes less generous.

3.4 MODEL FOR QUANTITATIVE ANALYSIS

While the stylised model of the previous section illustrated the relevant tradeoffs of supranational unemployment insurance, we would like to know whether our results survive in a more general setting. We therefore set up a dynamic model in the spirit of the simple model above, and include several additional features such as search effort and imperfect substitutability between Home and Foreign goods. We calibrate the model to Eurozone data and numerically solve for the Ramsey-optimal policies.

3.4.1 *Model setup*

Time is discrete at $t = 0, 1, 2, \ldots$. As before, a unit mass of workers and firms populates the economy, where $\omega \in (0, 1)$ workers live in the Home country and $(1 - \omega)$ workers live in the Foreign country. We describe the model setup in the Home country. The structure of the Foreign economy is identical to the Home country, but its parameter values (productivity levels, matching efficiency etc.) can be different.

3.4.1.1 *Workers*

A worker in the Home country can be employed or unemployed (indexed by $j = e, u$). It maximises expected lifetime utility

$$
\begin{aligned}
U_j &= \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \left(u\left(c_{jt}\right) - k\left(e_{jt}\right)\right)\right] \\
& c_{jt}, e_{jt} \geq 0 \\
u\left(c_{jt}\right) &= \frac{\left(c_{jt}\right)^{1-\gamma}}{1-\gamma}, \gamma \geq 0 \\
k\left(e_{jt}\right) &= \frac{\left(e_{jt}\right)^{1+\phi}}{1+\phi}, \phi > 0
\end{aligned}
$$

$\mathbb{E}[\cdot]$ is the expectations operator, $\beta \in (0,1)$ is the discount factor, e_{jt} is effort spent on job search (which is zero when $j = e$) and $k\left(e_{jt}\right)$ is the convex cost of job search. γ is the coefficient of relative risk aversion, and C_{jt} denotes expenditure on a consumption basket. This basket consists of goods produced in the Home and Foreign country and is given by

$$c_{jt} = \left(\psi\left(c_{jt,H}\right)^{\sigma} + (1-\psi)\left(c_{jt,F}\right)^{\sigma}\right)^{1/\sigma} \tag{3.39}$$

where $c_{jt,H}$ is the amount of goods consumed and produced at Home, while $c_{jt,F}$ is the amount of goods consumed at Home and produced in Foreign. The parameter $\sigma \in (-\infty, 1)$ governs the elasticity of substitution between foreign and domestic goods, which is constant at $1/(1-\sigma)$, and the parameter ψ represents the relative valuation of Home goods.[9]

We assume that there are no international trade costs, so the law of one price holds for both goods. We normalise the price of the Home good to one and denote with p_t be the price of Foreign goods. Thus, p_t equals the terms of trade of the Home country. Next, we define the *consumer price index* (CPI) at Home by $P_t = \left(c_{jt,H} + p_t c_{jt,F}\right)/c_{jt}$. Utility maximisation implies that

$$\frac{c_{jt,H}}{c_{jt,F}} = \left(p_t \frac{\psi}{1-\psi}\right)^{\frac{1}{1-\sigma}} \tag{3.40}$$

$$P_t = \left(\psi^{\frac{1}{1-\sigma}} + (1-\psi)^{\frac{1}{1-\sigma}} p_t^{-\frac{\sigma}{1-\sigma}}\right)^{-\frac{1-\sigma}{\sigma}} \tag{3.41}$$

We still need to specify workers' budget constraints and the financial assets they have access to. We want to capture an incomplete market setting in which workers can neither obtain perfect insurance

9 In the case of unitary elasticity of substitution ($\sigma = 0$), the consumption basket is of the Cobb-Douglas form $c_{jt} = \left(c_{jt,H}\right)^{\psi}\left(c_{jt,F}\right)^{1-\psi}$, so that the expenditure share on Home goods is exactly ψ. A situation where $\phi > \omega$ then corresponds to home bias in consumption.

of their idiosyncratic unemployment risk, nor perfect insurance of country-specific risk. As before, we will consider an extreme case in which workers do not have access to savings at all and simply consume their income each period. Intermediate forms of market incompleteness would certainly add realism, but at the cost of tractability of the model in the presence of heterogenous agents.

Employed workers receive the real wage w_t, of which an amount τ_t of payroll taxes is deducted. The unemployed receive unemployment insurance benefits b_t. Also, workers receive the profits from firms (described further below). We have to take a stance on how firm profits are distributed in the presence of heterogeneity, and assume that all agents own equal shares of the firms in their country which pay an aggregate profit π_t.[10] Thus, in each period the real value of an employed worker's consumption basket is simply his after-tax real wage plus profits, while the unemployed consume their unemployment benefit plus profits:

$$c_{et} = (w_t - \tau_t + \pi_t) / P_t \tag{3.42}$$

$$c_{ut} = (b_t + \pi_t) / P_t. \tag{3.43}$$

We are now ready to solve the agents' optimisation problem. Define the worker value functions as follows:

$$\mathcal{W}_t = u(c_{et}) + \beta \mathbb{E}_t [\mathcal{W}_{t+1} + s(\mathcal{U}_{t+1} - \mathcal{W}_{t+1})] \tag{3.44}$$

$$\begin{aligned} \mathcal{U}_t = \; & u(c_{ut}) + \beta \mathbb{E}_t [\mathcal{U}_{t+1} + f_{t+1} e_{t+1} (\mathcal{W}_{t+1} - \mathcal{U}_{t+1})] \\ & - \beta \mathbb{E}_t [k(e_{t+1})]. \end{aligned} \tag{3.45}$$

Then maximising the utility of the unemployed with respect to effort leads to the following optimality condition:

$$k'(e_t) = f_t (\mathcal{W}_t - \mathcal{U}_t). \tag{3.46}$$

3.4.1.2 *Firms*

Each country produces a distinct good. In the Home country, a representative firm produces the Home good using a production technology which is linear in labour:

$$y_t = a_t n_t. \tag{3.47}$$

Employment is subject to search frictions. The firm needs to post a number of vacancies v_t, each of which leads to successful matching with a worker with probability q_t. The vacancy filling rate is taken as given by the firm. Successful matches start production in the next period. At the same time, existing matches are destroyed at the exogenous rate s.

10 Firms will discount profits at a rate related to workers' marginal rates of substitution, as described below. Holding the firm portfolio therefore gives agents a limited form of savings through firms' intertemporal decisions, although they cannot save to insure their unemployment risk.

The firm needs to pay its workers a wage w_t (expressed in units of domestic goods), and it needs to pay a cost for each vacancy, which takes the form of a constant quantity of domestically produced goods κ_v. Its profits are given by

$$\pi_t = (a_t - w_t)\, n_t - \kappa_v v_t. \tag{3.48}$$

The firm maximises the discounted sum of profits

$$\mathbb{E} \sum_{t=0}^{\infty} Q_{0,t} \pi_t$$

where $Q_{s,t}$ is the discount factor between times s and t. Since the firm is owned in parts by employed and unemployed workers, it is not obvious what discount factor the firm should use. As in Jung and Kuester (2015), we set the firm discount factor to a weighted average of the worker discount factors:

$$Q_{s,t} = \beta^{t-s} \frac{n_t u'(c_{et}) + (1-n_t)\, u'(c_{ut})}{n_s u'(c_{es}) + (1-n_s)\, u'(c_{us})} \frac{P_s}{P_t}. \tag{3.49}$$

We denote by $\mathcal{J}_t$ the value of a filled job:

$$\mathcal{J}_t = a_t - w_t + (1-s)\, \mathbb{E}_t Q_{t,t+1} \left[\mathcal{J}_{t+1}\right]. \tag{3.50}$$

The optimality condition of the firm with respect to vacancy creation then takes the familiar form:

$$\kappa_v = q_t \mathcal{J}_t. \tag{3.51}$$

3.4.1.3 *Matching and wage determination*

At the beginning of period t, a fraction u_t of workers at Home are unemployed. We assume that labour is immobile across countries, so that workers can only search for jobs domestically. The number of total new hires is determined by the number of searching workers u_t, the search effort e_t of these workers, and the number of vacancies v_t. Workers and vacancies are then randomly matched according to a standard Cobb-Douglas matching function

$$m_t = \kappa_e \left(e_t u_t\right)^{\mu} v_t^{1-\mu} \tag{3.52}$$

where κ_e is a matching efficiency parameter and μ is the elasticity of matches with respect to unemployment. Defining labour market tightness as $\theta_t = v_t / e_t u_t$, the probability that a vacancy gets filled, and the probability that a worker putting in one unit of search effort finds a job, are given by:

$$q_t = m_t / v_t = \kappa^e \theta_t^{-\mu} \tag{3.53}$$

$$f_t = m_t / u_t = \kappa^e \theta_t^{1-\mu}. \tag{3.54}$$

Unemployed workers who separate have to wait one period before they can start searching again. Accordingly, the law of motion for employment and unemployment are given as follows:

$$n_t = (1-s)n_{t-1} + q_t v_t \tag{3.55}$$

$$u_t = 1 - n_{t-1}. \tag{3.56}$$

The wage paid to workers is determined by Nash bargaining in which workers and firms share the surplus from matching according to

$$\max_{w_t} \left(\mathcal{W}_t - \mathcal{U}_t\right)^{\xi} \mathcal{J}_t^{1-\xi}$$

where ξ is the bargaining power of workers. Due to the curvature of the utility function, a closed-form solution for the wage does not exist, but is implicitly given by the first-order condition:

$$\mathcal{W}_t - \mathcal{U}_t = \frac{\xi}{1-\xi} \frac{u'\left(c_{et}\right)}{P_t} \mathcal{J}_t. \tag{3.57}$$

3.4.1.4 *Government*

Unlike in the simple model of the previous section, we explicitly spell out national governments as well as a supranational unemployment insurance agency, each independently managing its finances.

The government in the Home country gains revenue exclusively from payroll taxes τ_{gt}. These taxes are used to fund benefits for unemployed workers b_{gt} as well as government expenditure g_t. Government expenditure is spent entirely on domestically produced goods.[11] The government has to balance its budget every period. Its budget constraint writes

$$g_t + u_t b_{gt} = \tau_{gt} n_t. \tag{3.58}$$

The supranational agency can likewise administer a component of unemployment insurance. This agency also has to balance its budget every period. It collects payroll taxes τ_{xt} and disburses unemployment benefits b_{xt} in the Home country, payroll taxes τ^*_{xt} and disburses unemployment benefits b^*_{xt} in the Foreign country. The agency's budget constraint writes

$$\omega\left(1-n_t\right) b_{xt} + (1-\omega)\left(1-n_t^*\right) p_t b_t^* = \omega n_t \tau_{xt} + (1-\omega) n_t^* p_t \tau_{xt}^*. \tag{3.59}$$

Total taxes on employed workers and total benefits received by unemployed workers, and the net replacement rate are then given by:

$$\tau_t = \tau_{gt} + \tau_{xt} \tag{3.60}$$

$$b_t = b_{gt} + b_{xt} \tag{3.61}$$

11 Our setup implicitly assumes that any utility workers receive from government expenditure is separable from market consumption, so that we can ignore it in the utility function.

In our benchmark calibration, the supranational agency is inactive ($b_{xt} = b^*_{xt} = \tau_{xt} = \tau^*_{xt} = 0$) and national governments target a constant replacement rate $\rho_t = \bar{\rho}$ and $\rho^*_t = \bar{\rho}^*$. Since this situation is close to the current system in place in the Eurozone, we refer to this as the "status quo".

3.4.1.5 *Market clearing and shocks*

The market clearing conditions for consumption goods produced in each country take the form:

$$\begin{aligned} \omega\left(y_t - \kappa_v v_t - g_t\right) &= \omega\left(n_t c_{et,H} + (1-n_t)\, c_{ut,H}\right) \\ &\quad + (1-\omega)\left(n^*_t c^*_{et,H} + (1-n^*_t)\, c^*_{ut,H}\right) && (3.62) \\ (1-\omega)\left(y^*_t - \kappa^*_v v^*_t - g^*_t\right) &= \omega\left(n_t c_{et,F} + (1-n_t)\, c_{ut,F}\right) \\ &\quad + (1-\omega)\left(n^*_t c^*_{et,F} + (1-n^*_t)\, c^*_{ut,F}\right). && (3.63) \end{aligned}$$

The exogenous shocks in our model are persistent shocks to productivity and government spending:

$$\log g_t = \rho_g \log g_{t-1} + \left(1-\rho_g\right) \log \bar{g} + \varepsilon_{gt} \quad (3.64)$$

$$\log a_t = \rho_a \log a_{t-1} + \left(1-\rho_a\right) \log \bar{a} + \varepsilon_{at}. \quad (3.65)$$

In particular, we rule out permanent shocks. This choice is not innocuous in our model, because it has implications for optimal risk sharing. The first best allocation in our model would completely shield domestic consumption from domestic employment and instead tie it to union output. In the presence of permanent shocks that differentially affect the long-run level of GDP in each country, this would effectively prescribe permanent fiscal transfers from the country with higher per capita income to the one with lower per capita income. We do not see much practical relevance and policital viability in such an extreme form of risk sharing. We therefore focus exclusively on mean-reverting shocks, so that cross-country transfers under the Ramsey planner will fall back to zero.

3.4.2 *Optimal policy*

Our goal is to characterise the optimal unemployment insurance policy of the government sector. We can consolidate the two national governments and the supranational agency by aggregating the budget constraints as in the simplified model:

$$\omega\left((1-n)\, b - n\tau\right) + (1-\omega)\left((1-n^*)\, pb^* - n^* p\tau^* -\right) = 0 \quad (3.66)$$

Effectively, the government sector has control over Home and Foreign benefits and taxes and faces one budget constraint, which implies three degrees of freedom. We express these degrees of freedom in terms of the Home and Foreign replacement rates $(\rho, \rho^*)$and the transfer from Foreign to Home as a fraction of Home GDP:

$$T_t = \frac{((1-n)\,b - n\tau)}{y_t} \tag{3.67}$$

A social planner who only faces the economy's resource constraints and search frictions in labour markets would simply equate marginal utilities of all workers in both countries (up to constant multiplying factors) and implement the efficient level of job creation as prescribed by the Hosios condition. But in our setup, there are not enough policy instruments to neutralise all three sources of inefficiency (undiversified idiosyncratic unemployment risk, search externalities and undiversified cross-country risk).

We therefore solve for the Ramsey-optimal policy where the planner maximises the objective function:

$$\mathbb{E}\left[\sum_{t=0}^{\infty}\beta^t\left(\begin{array}{c}\tilde{\omega}\,(n_t u\,(c_{et}) + (1-n_t)\,u\,(c_{ut})) \\ +\,(1-\tilde{\omega})\,(n_t^* u\,(c_{et}^*) + (1-n_t^*)\,u\,(c_{ut}^*))\end{array}\right)\right] \tag{3.68}$$

subject to all the equilibrium conditions of the economy. As in the simple model, we choose the planner weight on the Home country $\tilde{\omega}$ to rule out permanent transfers from one country to another:

$$\mathbb{E}\,[T_t] = 0. \tag{3.69}$$

3.4.3 *Calibration and model-data comparison*

We set the discount factor β to the standard value of 0.99 which yields an annual interest rate of 4 percent. The parameter σ is set to 0.736, implying an elasticity of substitution between Home and Foreign goods of 3.9 matching the European average of estimates reported in Corbo and Osbat (2013). Given that value, we calculate a value for the home good preference ψ of 0.56 to meet the corresponding average estimates of trade openness from Balta and Delgado (2009). The curvature of consumption γ is set to 1.5 as reported in Smets and Wouters (2003) and the search effort parameter is set to $\phi = 1$, corresponding to a unitary search effort elasticity. The effort scaling parameter κ_e is set to 0.692 to normalise steady state effort to unity.

We set the matching elasticity μ to the conventional value 0.5 according to estimates by Burda and Wyplosz (1994). The bargaining power of workers ξ is set lower, to 0.3. We target a steady-state unemployment rate of 9% and a quarterly vacancy-filling probability of 70% following Christoffel et al. (2009). A quarterly job finding rate of 30% is targeted in line with evidence provided by Elsby et al. (2013) for a number of European countries. The quarterly separation rate is then deduced from the implied steady state restrictions as $s = 0.030$. We also know that, in the steady state, the number of matches must

Table 3.1: Baseline calibration.

Parameter	Symbol	Value
Country size	ω	0.5
Preferences		
Discount factor	β	0.99
Risk aversion	γ	1.5
Preference Home/Foreign goods	σ	0.736
Relative valuation of Home goods	ψ	0.56
Inverse elasticity of search effort	ϕ	1
Effort cost scaling	κ_e	0.692
Labour market		
Matching elasticity	η	0.5
Matching efficiency	κ_e	0.692
Bargaining power	ξ	0.3
Separation rate	s	0.03
Vacancy costs	κ_v	0.711
Technology		
Steady-state level	$\bar{a}$	1
Autocorrelation	ρ_A	0.95
Std. dev.	σ_A	0.0069
Government spending		
Steady-state level	$\bar{g}$	0.182
Autocorrelation	ρ_G	0.79
Std. dev.	σ_G	0.0047
Policy		
Net replacement rate	ρ	0.65

be equal to the number of separations. This allows us to calculate the matching efficiency $\kappa_m = 0.458$ and vacancy posting costs $\kappa_v = 0.711$.

We set the technology shock persistence to $\rho_a = 0.95$ and its standard deviation such that the output's standard deviations obtained from our model matches the standard deviation of output in the data ($\sigma_a = 0.0069$). The government spending process is parameterised to match detrended government expenditure data as in Christoffel et al. ($\rho_g = 0.79$, $\sigma_g = 0.0047$).

Finally, the net replacement rate is set at $\rho = 65\%$ that is the average across EMU countries taking into account short and long run benefits, again following Christoffel et al.. The calibration is summarised in Table 3.1.

3.5 RESULTS

3.5.1 *Moments*

Table 3.2 reports several second moments of the calibrated model at the status quo policy and compares them to the data. Only the standard deviation of output is calibrated to match the data. While the persistence of fluctuations in the model matches the data relatively well, the model suffers from a counterfactually low unemployment volatility. This is of course a well-known problem (Shimer, 2005): With Nash bargaining, wages track movements in productivity too closely and the job creation rate is almost acyclical. Indeed, the real wage in the model is more volatile than in the data and almost perfectly correlated with output. We present an alternative specification with rigid wages further below.

Table 3.2: Second order moments in benchmark calibration.

Variable	std. dev.	rel. to real GDP	corr. with real GDP	1st autocorr.
Real GDP	0.87 [0.87]	1.00 [1.00]	1.00 [1.00]	0.72 [0.88]
Real wage	0.86 [0.57]	1.01 [0.66]	0.99 [0.36]	0.72 [0.80]
Unemployment	0.04 [0.40]	0.05 [0.47]	-0.49 [-0.86]	0.92 [0.95]

Second moments as obtained from simulating a linear approximation of the model at benchmark calibration. Real GDP is y_t/P_t, the real wage is w_t/P_t and unemployment is u_t. Corresponding moments in the data in parenthesis (from the ECB AWM database, 1984Q1-2008Q1). The second column reports the standard deviation, the third column reports the standard deviation relative to real GDP, the fourth column reports the cross-correlation with real GDP, and the last column reports the first order autocorrelation. Real GDP and real wage are in logarithms. All series are HP-filtered with smoothing parameter 1600.

We numerically calculate the Ramsey-optimal policy and report its cyclical stance in Table 3.3.

Table 3.3: Cyclicality of the optimal unemployment insurance policy.

	corr. with y_t/P_t	corr. with y_t^*/P_t^*
Optimal transfer/GDP T_t	-0.61	0.61
Optimal replacement rate ρ_t	-0.13	0.88
Optimal replacement rate ρ_t (no transfers)	0.91	0.15

Correlation coefficients of simulated model data, unfiltered.

The first row of the table shows the correlation of transfers T_t to the Home country (as a percentage of GDP) with Home and Foreign GDP, respectively. As expected, transfers correlate negatively with Home and positively with Foreign GDP, in order to insure local consumption from changes in local output. The second row reports the cyclicality of the replacement rate, which goes in the same direction: The replacement rate becomes more generous in a recession in which Home output falls relative to Foreign. When Home productivity is relatively low, the local benefit from efficient production is lower compared to that of increased insurance. However, when Foreign GDP falls, the replacement rate at Home drops. This increases Home output which is then transferred to Foreign.

In the third row, we constrain the Ramsey planner to not transfer any resources between countries, i.e. we impose $T_t = 0$. This corresponds to the optimal policy carried out by national unemployment insurance policies only, and effectively shuts down the international risk sharing dimension of policy design. Consistent with the intuition of the simple model, the replacement rate becomes less countercyclical in the absence of cross-country transfers. Here, this effect is so strong that the correlation with Home GDP actually turns positive: Nationally optimal policies feature a procyclical replacement rate, while the optimal supranational policy has countercyclical replacement rates.

3.5.2 *Impulse responses*

The effects of the Ramsey policy can be illustrated further by looking at impulse response functions. Figure 3.3 depicts a negative productivity shock.

Panel (a) shows the response of the policy instruments which reflects the correlations presented above. Under the status quo (solid black line), replacement rates are constant and transfers between countries are zero. Under the Ramsey-optimal policy (red dashed line), the Home country receives a transfer from Foreign which amounts to more than 0.2% of Home GDP on impact. At the same time, the Home replacement rate becomes more generous while the Foreign rate becomes less generous as Home productivity falls. By contrast, when we shut down international risk sharing by imposing zero transfers

Figure 3.3: Impulse responses, negative Home productivity shock.

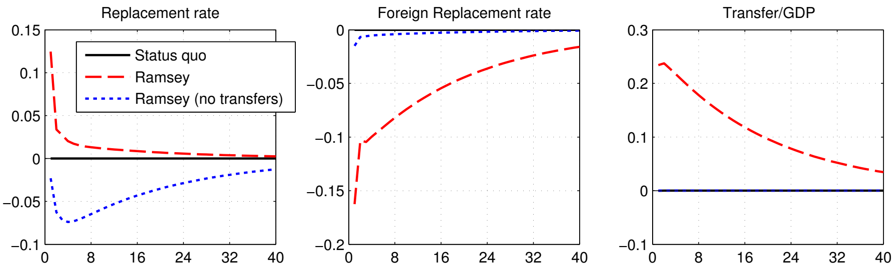

(a) Policy instruments.

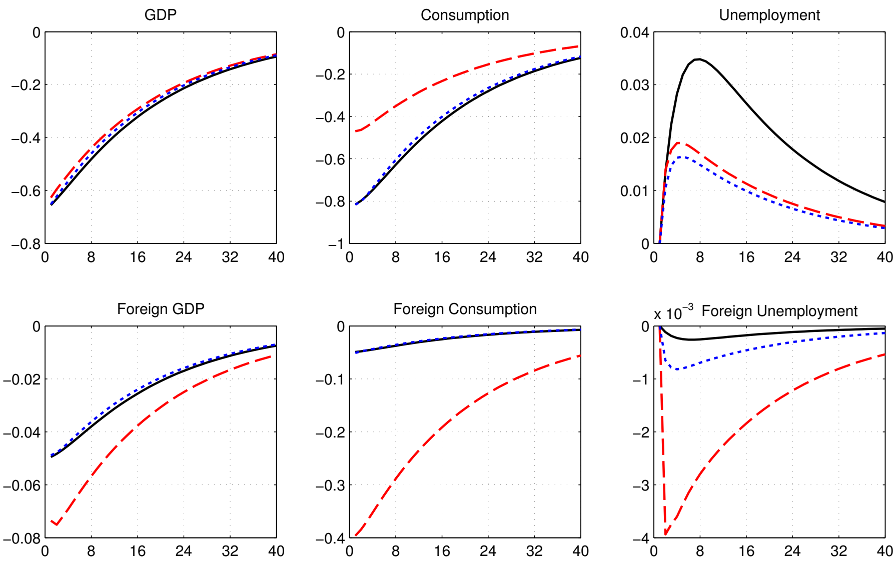

(b) Aggregate outcomes.

Impulse responses to a one-standard deviation negative Home productivity shock $\varepsilon_{At} = -\sigma_A = -0.69\%$. Replacement rates ρ_t, Foreign replacement rate ρ_t^*, Home transfer T_t , unemployment u_t and Foreign unemployment u_t^* are in percentage point deviation from steady-state. GDP y_t/P_t, Foreign GDP y_t^*/P_t^*, consumption c_t and foreign consumption c_t^* are in 100*log deviation from steady state.

(blue dotted line), the Home replacement rate falls instead (thereby becoming procyclical) while the Foreign rate barely reacts.

Panel (b) shows the effects of the shock on Home and Foreign output, consumption and the unemployment rate. The planner achieves a smaller rise in unemployment than under the status quo. This is not surprising since low bargaining power $\xi < \mu$ with Nash bargaining is known to lead to inefficiently volatile unemployment relative to the first-best. But the comparison to the Ramsey policy without transfers is revealing. Without transfers, the rise in unemployment is even lower. This is a reflection of the much lower replacement rate which increases search effort and reduces bargained wages. Foreign unemployment falls in all cases as the Foreign country experiences a positive terms of trade shock. Under the optimal policy, this fall is amplified by a lower replacement rate. In all cases however, the Foreign unemployment rate reacts very little.

The response of output is dominated by the fall in productivity and does not change much across policies. This also holds for consumption in the absence of transfers. But the optimal policy with transfers achieves a smaller reduction in consumption at Home but a sharper reduction in Foreign.

Figure 3.4 shows impulse responses to a positive shock to Home government spending. Under the status quo, the increase in government spending crowds out consumption and investment in vacancies, leading to a rise in unemployment and a fall in output under the status quo policy, as can be seen from Panel (b) of the figure. Consumption falls at Home and to a lesser extent in Foreign (since Foreign consumers also demand Home goods). The Ramsey-optimal policy reduces the replacement rate to mitigate the rise in unemployment, and effects a transfer from Foreign to Home to shield consumption from the crowding-out effect. For this shock, transfers are actually procyclical since real GDP rises at Home and falls in Foreign. This goes against our general finding of countercyclical transfers, but is a natural consequence of assuming that utility from government expenditure is separable from utility of market consumption, so that the social planner treats an increase in Home government spending like a pure loss of resources.

We end this section with a discussion of a policy proposal by Artus et al. (2013). They advocate a European unemployment insurance that pays unemployed workers a 20% net replacement rate, with national benefits reduced by the same amount so that the overall benefits are unaffected. This is financed by a contribution by the employed that is "set to 20% of the aggregate payroll multiplied by the structural unemployment rate in the country". The first thing to note about this proposal is that it does not run a balanced budget, and does not even take into account an intertemporal budget constraint. If we were to reproduce this policy in our model, even with perfectly known

Figure 3.4: Impulse responses, positive Home government spending shock.

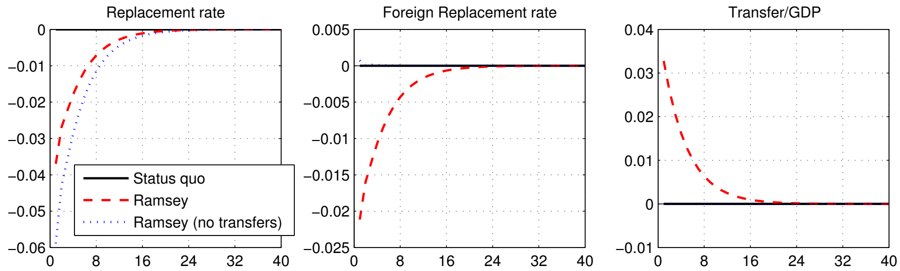

(a) Policy instruments.

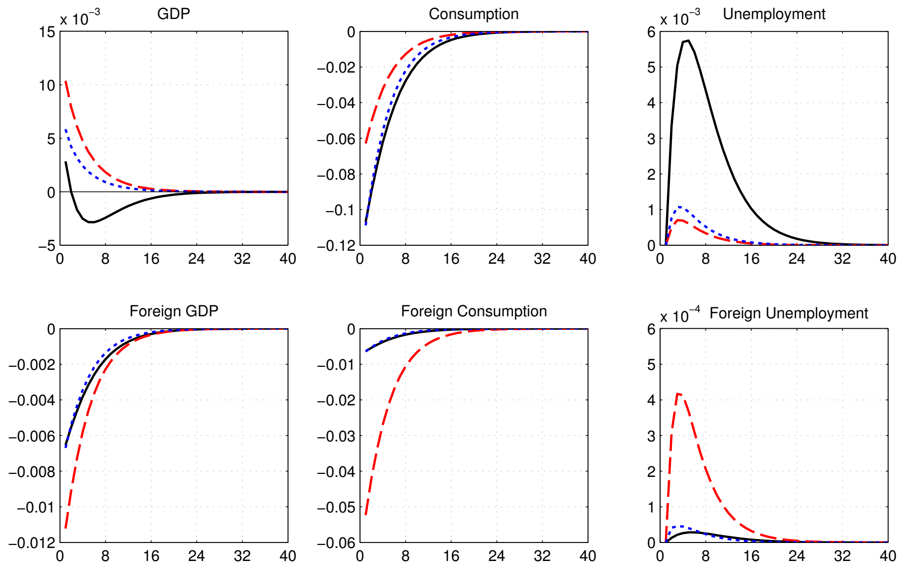

(b) Aggregate outcomes.

Impulse responses to a one-standard deviation positive Home government spending shock $\varepsilon_{Gt} = \sigma_G = 0.47\%$. Replacement rates ρ_t, Foreign replacement rate ρ_t^*, Home transfer T_t , unemployment u_t and Foreign unemployment u_t^* are in percentage point deviation from steady-state. GDP y_t/P_t, Foreign GDP y_t^*/P_t^*, consumption c_t and foreign consumption c_t^* are in 100*log deviation from steady state.

structural (steady-state) unemployment rates the financial position of the European unemployment scheme would have a unit root. It is not to us how Artus et al. conclude that the scheme would "avoid any permanent transfer" unless additional rules are put in place that determine how the deficits and surpluses are shared among member countries. In our model, we computed the effects of their proposal under the assumption that a constant fraction of the surplus or deficit of the scheme is born by each country, with a country's weight proportional to its steady-state share in union-wide GDP. We found that the effects were extremely small. However, this result is certainly due to the fact that the fluctuations in unemployment rates are counterfactually small. The Artus et al. proposal effectively bases the size of transfers on the difference between actual and structural unemployment rates, which is too small in our calibrated model to produce any appreciable effects. To address this problem, we now move on to a specification with rigid wages.

3.5.3 *Alternative specification with rigid wages*

One weakness of our benchmark calibration is the very low volatility of unemployment. In this section, we report results from an alternative specification of the model in which we make wages (measured in units of domestically produced goods) completely rigid by assuming

$$w_t = \bar{w}. \tag{3.70}$$

We keep the parameter values of our benchmark calibration as reported in Table 3.1, and choose the value of $\bar{w}$ such that the standard deviation of the unemployment rate matches exactly that in the data. We also adjust the standard deviation of the productivity shock σ_A to match the volatility of output in the data. This leads to a value of $\sigma_A = 0.0041$.

Table 3.4 compares the second moments to the data for our alternative specification. Unsurprisingly, the improvement in the behaviour

Table 3.4: Second order moments, rigid wage specification.

Variable	std. dev.	rel. to real GDP	corr. with real GDP	1st autocorr.
Real GDP	0.87 [0.87]	1.00 [1.00]	1.00 [1.00]	0.82 [0.88]
Real wage	0.07 [0.57]	0.09 [0.66]	-0.63 [0.36]	0.90 [0.80]
Unemployment	0.40 [0.40]	0.47 [0.47]	-0.71 [-0.86]	0.93 [0.95]

Second moments as obtained from simulating a linear approximation of the model at rigid wage specification. Corresponding moments in the data in parenthesis (from the ECB AWM database, 1984Q1-2008Q1). Notes from Figure 3.2 apply.

of the unemployment rate now comes at the expense of a counterfactually smooth wage rate.

Table 3.5: Cyclicality of the optimal unemployment insurance policy, rigid wage specification.

	corr. with y_t/P_t	corr. with y_t^*/P_t^*
Optimal replacement rate ρ_t	-0.45	-0.13
Optimal replacement rate ρ_t (no transfers)	-0.38	-0.01
Optimal transfer/GDP T_t	-0.46	0.46
Transfer/GDP T_t (Artus et al. proposal)	-0.66	0.64

Correlation coefficients of simulated model data, unfiltered.

Table 3.5 reports the cyclicalities of the optimal policy. One can see from the first row that our main result is unchanged: The Home replacement rate is countercyclical with respect to Home GDP. However, here it is also slightly negatively correlated with Foreign GDP. The optimal transfer (third row) also remains countercyclical. When we recalculate the optimal policy without transfers (second row), we find again that the Home replacement rate is less strongly correlated with output.

We then report the properties of the Artus et al. (2013) proposal as described above in the last row of Table 3.5. With realistic fluctuations in the unemployment rate, this policy proposal can in principle have sizeable effects. As expected, the transfer is countercyclical.

Here, we only report impulse response functions for a productivity shock, as this shock is the main driver of output fluctuations in the model. Figure 3.5 depicts the impulse responses to a negative Home productivity shock. As before, the solid black line is the response under the status quo policy (constant replacement rate, no transfers) and the dashed red line is the response under the Ramsey-optimal policy. We compare this policy to the Artus et al. proposal (green dotted line).

Panel (a) shows the optimal response of the policy instruments. As with flexible wages, the optimal Home replacement rate rises when Home productivity falls, and the Home country receives a transfer from Foreign. The magnitudes of the responses are much larger than with flexible wages: The Home replacement rate rises by 5 percentage points on impact, and remains 0.5 percentage point above its steady state for more than two years. The optimal transfer is also very large and amounts to more than 0.6% of Home GDP on impact with a slow decay.

Panel (b) reveals that the higher replacement rate only initially leads to a higher rise in unemployment, with the unemployment rate being lower than under the status quo after one year. The fall in real GDP is also mitigated. Foreign unemployment also rises, albeit by a much smaller amount, reflected in a small drop in Foreign GDP. The transfer of resources from Home to Foreign makes that Home

Figure 3.5: Impulse responses with rigid wages, negative Home productivity shock.

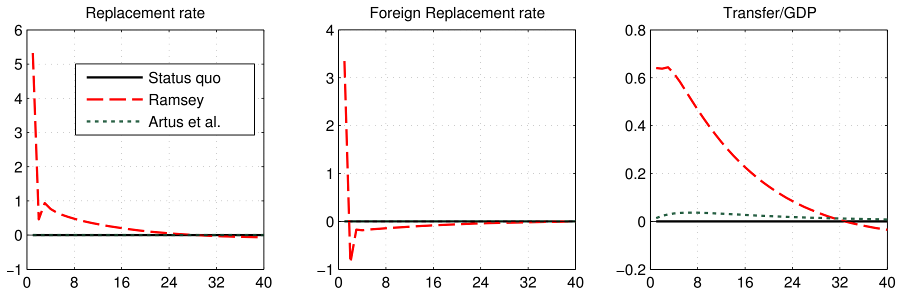

(a) Policy instruments.

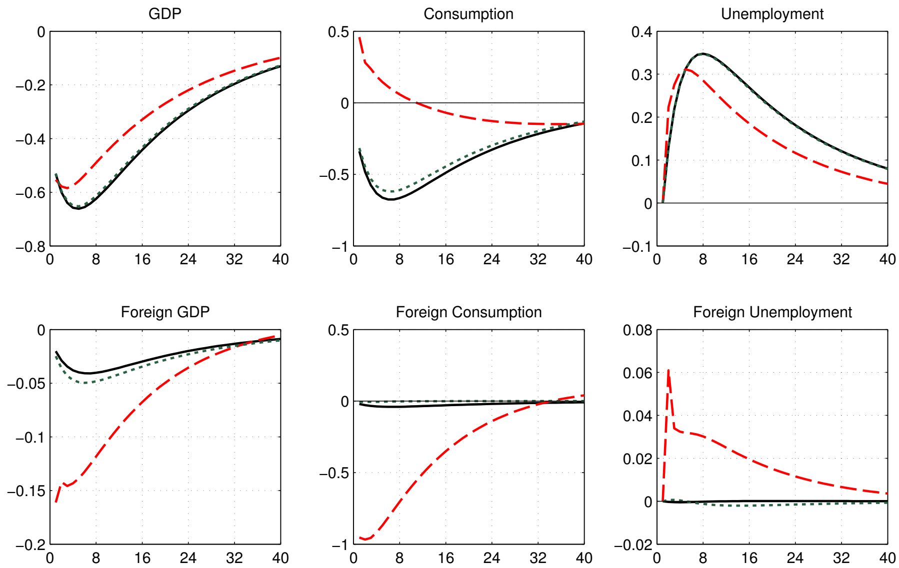

(b) Aggregate outcomes.

Impulse responses to a one-standard deviation negative Home productivity shock $\varepsilon_{At} = -\sigma_A = -0.69\%$. Replacement rates ρ_t, Foreign replacement rate ρ_t^*, Home transfer T_t , unemployment u_t and Foreign unemployment u_t^* are in percentage point deviation from steady-state. GDP y_t/P_t, Foreign GDP y_t^*/P_t^*, consumption c_t and foreign consumption c_t^* are in 100*log deviation from steady state.

consumption does not drop as much and even rises on impact, while Foreign consumption falls significantly.

The impact of the Artus et al. (2013) proposal (green dashed line in the figure) remains very limited even with unemployment fluctuations as large as in the data. The reduction in Home consumption, for example, is somewhat smaller than under the status quo policy, but the difference is very small. In light of this result, the benefit of this proposal seems questionable.

3.6 CONCLUSION

In this chapter, we used an international business cycle model augmented by frictional labour markets and incomplete financial markets to discuss optimal unemployment insurance policy operating across multiple countries. This adds international risk sharing to the classic policy tradeoff between efficient employment and insurance of unemployment risk. We have shown that cross-country insurance through the unemployment insurance system can in principle be achieved without affecting unemployment levels; and that the desirability of international risk-sharing introduces a countercyclical element to the optimal unemployment insurance policy. Calibrated to Eurozone data, our model implied that the international risk-sharing component dominates in the design of optimal policy, making it countercyclical overall. The optimal policy prescribes significant transfers between countries as well as countercyclical replacement rates. By contrast, recent policy proposals seem to have only a limited impact on business cycle dynamics and international risk sharing.

There are several directions in which our findings could be extended. First, we currently employ a very stylised model of the Eurozone economy, with symmetrical countries, no private or public savings possibilities and no nominal rigidities, thereby abstracting from many potentially relevant factors for the optimal policy design. In these dimensions, our analysis can be refined further. Second, the optimal policy we compute here is one in which the planner has perfect knowledge of the structure of the economy. One of the most difficult issues in implementing a policy such as the one in this chapter is that the structural rate of unemployment can only be reliably estimated in hindsight, if at all. It would be useful to see whether simple policy rules that are more easily implementable under imperfect information can reasonably approximate the optimal policy.

A APPENDIX

A.1 SECOND-ORDER PERTURBATION METHOD FOR CHAPTER 1

It is straightforward to apply perturbation techniques (see for example Schmitt-Grohe and Uribe 2004) to solve for an internally rational equilibrium with conditionally model-consistent expectations. To summarise the equilibrium conditions, the actual policy function g is

$$g\left(y,u,\sigma\right)=h\left(y,u,r\left(y,u,\sigma\right),\sigma\right) \tag{A.1}$$

where the subjective policy function h solves

$$\begin{aligned} 0=\Psi\left(y,u,z,\sigma\right) &= \begin{pmatrix} \psi\left(y_t,y_{t-1},u_t,z_t\right) \\ \mathbb{E}_t^{\mathcal{P}}\left[f_{-j_0}\left(y_{t+1},y_t,y_{t-1},u_t\right)\right] \end{pmatrix} \\ =\int_{\mathbb{R}^p}\int_{\mathbb{R}} &\begin{pmatrix} \psi\left(h\left(y,u,z,\sigma\right),y,u,z\right) \\ f_{-j_0}\left(h\left(h\left(y,u,z,\sigma\right),u',z',\sigma\right),h\left(y,u,z,\sigma\right),y,u\right) \end{pmatrix} G\left(dz'\right)F\left(du'\right) \end{aligned} \tag{A.2}$$

and the equilibrium subjective forecast error function r solves

$$\begin{aligned} 0 &= \mathbb{E}_t^{\mathcal{P}}\phi\left(y_{t+1},y_t,y_{t-1},u_t\right) \\ &= \int_{\mathbb{R}^p}\int_{\mathbb{R}}\phi\left(h\left(h\left(y,u,r\left(y,u,\sigma\right),\sigma\right),u',z',\sigma\right),h\left(y,u,r\left(y,u,\sigma\right),\sigma\right),y,u\right)G\left(dz'\right)F\left(du'\right) \\ &= \Phi\left(y,u,\sigma\right) \end{aligned} \tag{A.3}$$

The goal is to derive an accurate second-order approximation of the objective policy function g around the non-stochastic steady state $\bar{y}$:

$$\begin{aligned} g\left(y,u,\sigma\right) &\approx g\left(\bar{y},0,0\right) \\ &+ g_y\left(y-\bar{y}\right)+g_u u+g_\sigma\sigma \\ &+ \frac{1}{2}g_{yy}\left[\left(y-\bar{y}\right)\otimes\left(y-\bar{y}\right)\right]+\frac{1}{2}g_{yu}\left[\left(y-\bar{y}\right)\otimes u\right]+\frac{1}{2}g_{uu}\left[u\otimes u\right] \\ &+ \frac{1}{2}g_{\sigma\sigma}\sigma^2 \end{aligned} \tag{A.4}$$

The first step in deriving the approximation consists in calculating this approximation for the subjective policy function h. This can be done using standard methods. The second step consists in finding the derivatives of the function r. Total differentiation at the non-stochastic steady state leads to the following first-order derivatives:

$$0 = \frac{d\Phi}{dy}(\bar{y},0,0) \quad = \quad \left(\phi_{y_{t+1}} h_y + \phi_{y_t}\right)\left(h_y + h_z r_y\right) + \phi_{y_{t-1}} \quad \text{(A.5)}$$

$$0 = \frac{d\Phi}{du}(\bar{y},0,0) \quad = \quad \left(\phi_{y_{t+1}} h_y + \phi_{y_t}\right)\left(h_u + h_z r_u\right) + \phi_{u_t} \quad \text{(A.6)}$$

$$0 = \frac{d\Phi}{d\sigma}(\bar{y},0,0) \quad = \quad \left(\phi_{y_{t+1}} h_y + \phi_{y_t}\right)\left(h_\sigma + h_z r_\sigma\right) \quad \text{(A.7)}$$

I assume that the equilibrium conditions imply that $\bar{z} = 0$ at the steady-state. This means that in the absence of shocks, agents make no forecast errors under learning.

Define the matrix $A = \left(\phi_{y_{t+1}} h_y + \phi_{y_t}\right) h_z$. Then the first-order derivatives of r are given by:

$$r_y \quad = \quad -A^{-1}\left(\left(\phi_{y_{t+1}} h_y + \phi_{y_t}\right) h_x + \phi_{y_{t-1}}\right) \quad \text{(A.8)}$$

$$r_u \quad = \quad -A^{-1}\left(\left(\phi_{y_{t+1}} h_y + \phi_{y_t}\right) h_u + \phi_{u_t}\right) \quad \text{(A.9)}$$

$$r_\sigma \quad = \quad 0 \quad \text{(A.10)}$$

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

The second-order calculations are similar. The second-order derivative of Φ with respect to y is:

$$\begin{aligned} 0 = \frac{d^2\Phi}{dy^2}(\bar{y},0,0) \quad &= \quad \left(\phi_{y_{t+1}} g_y + \phi_{y_t}\right)\left(h_{yy} + 2h_{yz}\left[I_n \otimes r_y\right] + h_{zz}\left[r_y \otimes r_y\right]\right) \\ &+ \quad \phi_{y_{t+1}} h_{yy}\left[g_y \otimes g_y\right] + B_{yy} + A r_{yy} \qquad \text{(A.11)} \end{aligned}$$

This is a system of n^2 linear equations in r_{yy}, and thus can be solved easily. Again, only invertibility of the matrix A is required for a unique solution. The cross-derivatives of ϕ are collected in the matrix B_{yy} (of size $n_z \times n^2$), which contains only first-order derivatives of the policy functions:

$$\begin{aligned} B_{yy} \quad = \quad & \phi_{y_{t+1}y_{t+1}}\left[h_y g_y \otimes h_y g_y\right] + \phi_{y_t y_t}\left[g_y \otimes g_y\right] + \phi_{y_{t-1}y_{t-1}} + \phi_{z_t z_t}\left[r_y \otimes r_y\right] \\ + \quad & 2\phi_{y_{t+1}y_t}\left[h_y g_y \otimes g_y\right] + 2\phi_{y_{t+1}y_{t-1}}\left[h_y g_y \otimes I_n\right] + 2\phi_{y_{t+1}z_t}\left[h_y g_y \otimes r_y\right] \\ + \quad & 2\phi_{y_t y_{t-1}}\left[g_y \otimes I_n\right] + 2\phi_{y_t z_t}\left[g_y \otimes r_y\right] + 2\phi_{y_{t-1}z_t}\left[I_n \otimes r_y\right] \qquad \text{(A.12)} \end{aligned}$$

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

$$0 = \frac{d^2\Phi}{dyd\sigma}(\bar{y},0,0) = \left(\phi_{y_{t+1}} g_y + \phi_{y_t}\right)\left(h_{y\sigma} + h_z r_{y\sigma}\right) + \phi_{z_t} r_{y\sigma} \quad \text{(A.13)}$$

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

$$\begin{aligned} 0 = \frac{d^2\Phi}{d\sigma^2}(\bar{y}, 0, 0) &= \phi_{y_{t+1}y_{t+1}} \left(h_{\sigma\sigma} + h_{uu}\text{vec}\left(\Sigma_u\right) + +h_{zz}\text{vec}\left(\Sigma_z\right)\right) \\ &+ \phi_{y_{t+1}y_{t+1}} \left(\text{vec}\left(h_u' \Sigma_u h_u\right) + \text{vec}\left(h_z' \Sigma_z h_z\right)\right) \\ &+ \left(\phi_{y_{t+1}} g_y + \phi_{y_t}\right)\left(h_{\sigma\sigma}\right) + A r_{\sigma\sigma} \end{aligned} \tag{A.14}$$

Again, this can be solved for $r_{\sigma\sigma}$ when A is invertible. Note that the perceived variance Σ_z appears in the calculation because it matters for expectations for the future (unless $\phi_{y_{t+1}y_{t+1}} = 0$).

The second-order derivatives of the actual policy function g are calculated easily once those of r are known using the Chain rule. In particular, the cross-derivatives $g_{u\sigma}$ and $g_{y\sigma}$ are zero, just as under rational expectations.

A.2 FURTHER DETAILS ON THE FULL MODEL IN CHAPTER 2

Here, I provide further details on the model presented in Section 2.5 of Chapter 2.

Retailers

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

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

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

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

$$\frac{P_t^*}{P_t} = \frac{1}{1+\tau}\frac{\sigma}{\sigma-1}\frac{\Gamma_{1t}}{\Gamma_{2t}} \tag{A.15}$$

$$\Gamma_{1t} = q_t + \kappa \mathbb{E}_t^{\mathcal{P}} \Lambda_{t+1} \frac{\tilde{Y}_{t+1}}{\tilde{Y}_t} \pi_{t+1}^{\sigma} \tag{A.16}$$

$$\Gamma_{2t} = 1 + \kappa \mathbb{E}_t^{\mathcal{P}} \Lambda_{t+1} \frac{\tilde{Y}_{t+1}}{\tilde{Y}_t} \pi_{t+1}^{\sigma-1} \tag{A.17}$$

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

$$1 = (1-\kappa)\left(\frac{P_t^*}{P_t}\right)^{1-\sigma} + \kappa \pi_t^{\sigma-1} \tag{A.18}$$

and the Tak-Yun distortion term is

$$\Delta_t = (1-\kappa)\left(\frac{\Gamma_{1t}}{\Gamma_{2t}}\right)^{-\sigma} + \kappa \pi_t^{\sigma} \Delta_{t-1}. \tag{A.19}$$

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.

Labour agencies

Similarly to retailers, labour agencies transform the homogeneous household labour input into differentiated labour 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. Labour agency h solves the following optimisation:

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

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

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

$$w_t = \frac{W_t}{P_t} = w_{t-1}\frac{\pi_{wt}}{\pi_t}. \tag{A.20}$$

Capital good producers

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

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

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

$$Q_t = 1 + \psi \left(\frac{I_t}{I_{t-1}} - 1 \right) \tag{A.21}$$

Market clearing

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

$$Y_t = \int_0^1 Y_{jt} dj = \int_0^1 Y_{it} di \tag{A.22}$$

$$\tilde{Y}_t = \frac{Y_t}{\Delta_t} = C_t + I_t + -\frac{\psi}{2} \left(\frac{I_t}{I_{t-1}} - 1 \right)^2 + C_t^e \tag{A.23}$$

$$L_t = \int_0^1 L_{ht} dh \tag{A.24}$$

$$\tilde{L}_t = \frac{L_t}{\Delta_{wt}} = \int_0^1 L_{jt} dj \tag{A.25}$$

$$K_t = \int_0^1 K_{jt} dj = (1 - \delta) K_{t-1} + I_t \tag{A.26}$$

$$1 = S_{jt},\ j \in [0, 1] \tag{A.27}$$

$$0 = B_t^g \tag{A.28}$$

A.3 PROPERTIES OF THE RATIONAL EXPECTATIONS EQUILIBRIUM IN CHAPTER 2

The rational expectations equilibrium considered in Section 2.5 has the following properties that need to be verified. All statements are local in the sense that for each of them, there exists a neighbourhood of the non-stochastic steady-state in which the statement holds.

1. All firms choose the same capital-labour ratio K_{jt}/L_{jt} .

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

3. At any time t, the stock market valuation P_{jt} of a firm j is proportional to its net worth after entry and exit $\tilde{N}_{jt}$ with a slope that is strictly greater than one.
4. Borrowers never default on the equilibrium path and borrow at the risk-free rate, and the lender only accepts debt payments up to a certain limit.
5. If the firm defaults and the lender seizes the firm, she always prefers restructuring to liquidation.
6. The firm always exhausts the borrowing limit.
7. All firms can be aggregated. Aggregate debt, capital and net worth are sufficient to describe the intermediate goods sector.

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

Value functions

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

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

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

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

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

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

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

Characterising the first stage

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

$$R^k = \alpha q \left((1-\alpha)\frac{qA}{w}\right)^{\frac{1-\alpha}{\alpha}} + (1-\delta) Q \qquad \text{(A.29)}$$

Now taking Property 3 as given for now, Y_2 is a linear function with slope strictly greater than one. Then the following holds for the first-stage value function Y_1:

$$\begin{aligned} Y_1(K,B) &= N + (1-\gamma)(D - N + Y_2(N-D)) \\ &= N + (1-\gamma)\left(Y_2' - 1\right)((1-\zeta)N + \zeta(QK - B)) \\ &> N \\ &= R^k K - RB \end{aligned} \qquad \text{(A.30)}$$

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

Characterising the second stage

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

1. The firm (F) proposes a borrowing contract $\left(B, \tilde{R}\right)$.
2. The lender (L) can accept or reject the contract.
   - A rejection corresponds to setting the contract $\left(B, \tilde{R}\right) = (0, 0)$.
     Payoff for L: 0. Payoff for F: $\tilde{\beta}\mathbb{E}\left[Y_1\left(\tilde{N}, 0\right)\right]$.
3. F acquires capital and can then choose to default or not.
   - If F does not default, it has to repay in the next period.
     Payoff for L: $\mathbb{E}Q_{t,t+1}\tilde{R}B - B$. Payoff for F: $\tilde{\beta}\mathbb{E}\left[Y_1\left(K, \frac{\tilde{R}}{R}B\right)\right]$.
4. If F defaults, the debt needs to be renegotiated. F makes an offer for a new debt level B^*.[1]
5. L can accept or reject the offer.

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

- If L accepts, the new debt level replaces the old one.
  Payoff for L: $\mathbb{E}\Lambda\tilde{R}B^* - B$. Payoff for F: $\tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(K, \frac{\tilde{R}}{R}B^*\right)\right]$.

6. If L rejects, then she seizes the firm. A fraction $1-\xi$ of the firm's capital is lost in the process. Nature decides randomly whether the firm can be "restructured".
   - If the firm cannot be restructured, or it can but the lender chooses not to do so, then the lender has to liquidate the firm.
     Payoff for L: $\mathbb{E}\Lambda\xi QK - B$. Payoff for F: 0.
   - If the firm can be restructured and the lender chooses to do so, she retains a debt claim of present value ξB and sells the residual equity claim in the firm to another investor.
     Payoff for L: $\xi B + \tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(\xi K, \xi B\right)\right] - B$. Payoff for F: 0.

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

$$\xi B + \tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(\xi K, \xi B\right)\right] \geq \mathbb{E}\Lambda\xi QK. \tag{A.31}$$

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

$$\begin{aligned}\xi B + \tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(\xi K, \xi B\right)\right] &> \xi B + \tilde{\beta}\mathbb{E}\left[R^k\xi K - R\xi B\right] \\ &= \tilde{\beta}\mathbb{E}\left[R^k\xi K\right] \\ &> \xi K \end{aligned} \tag{A.32}$$

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

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

$$\mathbb{E}\Lambda\tilde{R}B^* \geq x\left(\xi B + \tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(\xi K, \xi B\right)\right]\right) + (1-x)\,\mathbb{E}\Lambda\xi QK. \tag{A.33}$$

Now turn to the firm F. Among the set of offers B^* that are accepted by L, the firm will prefer the lowest one, i.e. that which satisfies (A.33) with equality. This follows from Υ_1 being a decreasing function of debt. This lowest offer will be made if it leads to a higher payoff than expropriation: $\tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(K, \frac{\tilde{R}}{R}B^*\right)\right] \geq 0$. Otherwise, F offers zero and L seizes the firm.

Going one more step backwards, F has to decide whether to declare default or not. It is preferable to do so if the B^* that L will just accept is strictly smaller than B or if expropriation is better than repaying, $\tilde{\beta}\mathbb{E}\left[Y_1\left(K, \frac{\tilde{R}}{R}B\right)\right] \geq 0$.

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

$$\begin{aligned} \tilde{\beta}\mathbb{E}\left[Y_1\left(K,B\right)\right] &> \tilde{\beta}\mathbb{E}\left[R^k K - RB\right] \\ &= \tilde{\beta}\mathbb{E}\left[R^k \tilde{N} + \left(R^k - R\right)B\right] \\ &> 0. \end{aligned} \tag{A.34}$$

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

$$\max_{\tilde{R},B,B^*} \tilde{\beta}\mathbb{E}\left[Y_1\left(\tilde{N} + B, \frac{\tilde{R}}{R}B^*\right)\right]$$

$$\begin{aligned} \text{s.t. } \mathbb{E}\Lambda\tilde{R}B^* &\geq B \\ \mathbb{E}\Lambda\tilde{R}B^* &= x\left(\xi B + \tilde{\beta}\mathbb{E}\left[Y_1\left(\xi\left(\tilde{N}+B\right), \xi B\right)\right]\right) \\ &\quad + (1-x)\,\mathbb{E}\Lambda Q\xi\left(\tilde{N}+B\right) \end{aligned}$$

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

$$\frac{\xi}{R}\left(1 - x + xR + xY_1'\left[\frac{R^k}{R} - 1\right]\right) < 1 \tag{A.35}$$

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

Linearity of firm value

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

$$
\begin{aligned}
\Upsilon_2\left(\tilde{N}\right) &= \tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(\tilde{N}+B,B\right)\right] && \text{(A.36)}\\
\text{where } B &= x\left(\xi B+\tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(\xi\left(\tilde{N}+B\right),\xi B\right)\right]\right)\\
&\quad+\left(1-x\right)Q_t\xi\left(\tilde{N}+B\right) && \text{(A.37)}
\end{aligned}
$$

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

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

$$
\begin{aligned}
\Upsilon_2' &= \tilde{\beta}\frac{\Upsilon_1\left(K,B\right)}{\tilde{N}}\\
&= \tilde{\beta}\frac{\gamma\left(R^kK-RB\right)+\left(1-\gamma\right)\Upsilon_2\left(R^kK-RB\right)}{\tilde{N}}\\
&= \tilde{\beta}\left(\gamma+\left(1-\gamma\right)\Upsilon_2'\right)\left(R^k\frac{K}{\tilde{N}}-R\frac{B}{\tilde{N}}\right)\\
&= \left(\gamma+\left(1-\gamma\right)\Upsilon_2'\right)\underbrace{\frac{R^k+\left(R^k-R\right)\frac{B}{\tilde{N}}}{R}}_{=:c_0>1}\\
&= \frac{\gamma c_0}{1-\left(1-\gamma\right)c_0}\\
&> 1 && \text{(A.38)}
\end{aligned}
$$

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

$$
\begin{aligned}
K_t=\int_0^1 K_{jt}dj &= \int_{j\notin\Gamma_t}\left(N_{jt}-\zeta E_{jt}+B_{jt}\right)dj+\int_{j\in\Gamma_t}\left(\omega\left(N_t-\zeta E_t\right)+B_{jt}\right)dj\\
&= \left(1-\gamma+\gamma\omega\right)\left(N_t-\zeta E_t\right)+B_t && \text{(A.39)}\\
N_t=\int_0^1 N_{jt}dj &= R_t^kK_{t-1}-R_{t-1}B_{t-1} && \text{(A.40)}\\
B_t=\int_0^1 B_{jt}dj &= x\xi\left(B_t+P_t\right)+\left(1-x\right)\xi\mathbb{E}_t\Lambda_{t+1}Q_{t+1}K_t && \text{(A.41)}
\end{aligned}
$$

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

Return on capital

It can now be shown under which conditions the internal rate of return is indeed greater than the return on debt. From the steady-state versions of equations (A.39) and (A.40), it follows that

$$
R^k=R+\left(G-R\left(1-\gamma+\gamma\omega\right)\right)\frac{\bar{N}}{\bar{K}}+Rc\left(1-\gamma+\gamma\omega\right)\frac{\bar{E}}{\bar{K}}. \qquad \text{(A.42)}
$$

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

$$\gamma > \frac{R - G}{G(1-\omega)}. \tag{A.43}$$

A.4 PROPOSITIONS FOR CHAPTER 3

Here, we provide the mathematical results for Section 3.3.

The optimal replacement rate in the absence of private risk sharing satisfies Equation (3.37) in the main text:

$$\underbrace{\frac{(1-n)(1-\rho)}{n+(1-n)\rho} - \epsilon_\rho^n \left(\log\rho + \frac{1-\rho}{n+(1-n)\rho}\right)}_{=:I(\rho)} = \underbrace{-\epsilon_\rho^y \frac{1}{n} x}_{=:H(\rho)}$$

where $x = y/c$. Throughout, we make the following assumption:

Assumption A.1. *$I(\rho)$ is strictly concave, $H(\rho)$ and $(H \cdot y)(\rho)$ are strictly convex in ρ on $[0,1]$.*

We numerically verified Assumption 3 for a wide range of parameters, and conjecture that Assumption (A.1) always holds true. The limit behaviour of the functions at the corners is easy to prove and together with our assumption determines the shape of the curves in the main text.

Proposition A.2. *$I(0) = \frac{1-\bar{n}}{\bar{n}}$ where $\bar{n} = \kappa_m \left(\frac{\kappa_m}{\kappa_v} a\right)^{(1-\mu)/\mu}$ and $I(1) = 0$. Also, holding y/c constant, $H(0) = H\left(exp\left(-\frac{1-\mu}{\mu}\frac{\xi}{1-\xi}\right)\right) = 0$, $\lim_{\rho\to 1} = +\infty$, and $H'(\rho)$ is strongly convex in $[0,1]$.*

Proof. We start with the insurance term $I(\rho)$. At the limit when $\rho \to 0$, we have $w \to 0$ and $n = \kappa_m \left(\frac{\kappa_m}{\kappa_v}(a-w)\right)^{(1-\mu)/\mu} \to \bar{n}$. Therefore:

$$\frac{(1-n)(1-\rho)}{n+(1-n)\rho} \xrightarrow{\rho\to 0} \frac{1-\bar{n}}{\bar{n}}. \tag{A.44}$$

The remaining term of $I(\rho)$ must therefore go to zero. Indeed,

$$\begin{aligned} \epsilon_\rho^n = \frac{dn}{d\rho}\frac{\rho}{n} &= -\frac{a}{a-w}\frac{1-\mu}{\mu}\frac{w^2}{a^2}\frac{1-\xi}{\xi} \\ &= \frac{1}{\log\rho}\frac{1-\mu}{\mu}\frac{\xi}{\xi-(1-\xi)\log\rho} \end{aligned} \tag{A.45}$$

and therefore

$$\begin{aligned} &\epsilon_\rho^n \left(\log\rho + \frac{1-\rho}{n+(1-n)\rho}\right) \\ =& \frac{1-\mu}{\mu}\frac{\xi}{\xi-(1-\xi)\log\rho}\left(1 + \frac{1-\rho}{n+(1-n)\rho}\frac{1}{\log\rho}\right) \xrightarrow{\rho\to 0} 0. \end{aligned} \tag{A.46}$$

For the case $\rho \to 1$, the first term clearly disappears:

$$\frac{(1-n)(1-\rho)}{n+(1-n)\rho} \xrightarrow{\rho \to 1} 0 \tag{A.47}$$

and for the second term, we have:

$$\begin{aligned} \frac{1-\mu}{\mu} \frac{\xi}{\xi-(1-\xi)\log\rho} \left(1+\frac{1-\rho}{n+(1-n)\rho}\frac{1}{\log\rho}\right) \\ \xrightarrow{\rho \to 0} \frac{1-\mu}{\mu}\left(1+\lim_{\rho\to 1}\frac{1-\rho}{\log\rho}\right) &= 0. \end{aligned} \tag{A.48}$$

Next, we turn to the $H(\rho)$ function. As $\rho \to 0$, $n \to \bar{n} > 0$ and and $w \to 0$. Therefore

$$-\epsilon_\rho^y \frac{1}{n} = -\frac{1}{n}\frac{w}{a}\left(\frac{1-\xi}{\xi}+\frac{1-\mu}{\mu}\frac{1}{\log\rho}\right) \xrightarrow{\rho\to 0} 0. \tag{A.49}$$

And as $\rho \to 1$, $w \to 1$ and $n \to 0+$, so that

$$-\frac{1}{n}\frac{w}{a}\left(\frac{1-\xi}{\xi}+\frac{1-\mu}{\mu}\frac{1}{\log\rho}\right) \xrightarrow{\rho\to 1} +\infty. \tag{A.50}$$

□

Proposition A.3. *The optimal replacement rate is unique and strictly between* $\exp\left(\frac{\mu}{1-\mu}\frac{1-\xi}{\xi}\right)$ *and one.*

Proof. Since $f(\rho) = H(\rho) - I(\rho)$ is continuous on $[0,1]$ and a strictly concave by Assumption (A.1), it crosses zero at most twice. But $f(0) > 0$ and $\lim_{\rho\to} -\infty$, so there is a unique interior solution ρ^* to $f(\rho) = 0$. Since $I(0) > I(1) = 0$ and I is strictly concave, $I(\rho) > 0 \forall \rho \in (0,1)$ and the optimum has $H(\rho^*) > 0$. Since H is a strictly convex function, $H(0) = 0$ and $\lim_{\rho\to 1} H(\rho) = +\infty$ and , $H(\rho_0) = 0$ for exactly one $\rho_0 \in (0,1)$ and $\rho^* > \rho_0$. Finally, $H\left(\exp\left(\frac{\mu}{1-\mu}\frac{1-\xi}{\xi}\right)\right) = 0$. □

Proposition A.4. *The optimal replacement rate is strictly decreasing in* y/c.

Proof. Define $x = y/c$. Taking the total derivative of the optimality condition with respect to x, we have

$$\begin{aligned} 0 &= \frac{\partial I}{\partial x} - \frac{\partial H}{\partial x} + \frac{\partial I}{\partial \rho}\frac{d\rho}{dx} - \frac{\partial H}{\partial \rho}\frac{d\rho}{dx} \\ \Leftrightarrow \frac{d\rho}{dx} &= -\frac{\frac{\partial I}{\partial x} - \frac{\partial H}{\partial x}}{\frac{\partial I}{\partial \rho} - \frac{\partial H}{\partial \rho}}. \end{aligned} \tag{A.51}$$

Clearly, $dI/dx = 0$ and at the optimal ρ, we have $dH/dx = H(\rho)/x > 0$. Furthermore, we know that $I(0) > H(0)$ and $I(\rho) = H(\rho)$ only once, so it must be the case that $dH/d\rho > dI/d\rho$ at the optimal ρ. Therefore $d\rho/dx < 0$. □

Proposition A.5. *In the limit as $\omega \to 0$, the optimal replacement rate is unique, strictly below one and strictly decreasing in a as y/c is chosen optimally.*

Proof. As $\rho \to 0$, the risk-sharing condition (3.37) becomes

$$c = \frac{\mathbb{E}[y]}{\mathbb{E}[y^*]} y^*. \tag{A.52}$$

The optimal choice of ρ when y/c is chosen optimally can now be described as

$$I(\rho) = \tilde{H}(\rho) \tag{A.53}$$

$$\text{where } \tilde{H}(\rho) = H(\rho) \frac{\mathbb{E}[y^*]}{\mathbb{E}[y]} \frac{y}{y^*}.$$

By Assumption (A.1), $\tilde{H}$ is a strictly convex function. The behaviour of $\tilde{H}$ at zero is

$$\tilde{H}(0) = H(0) \frac{\mathbb{E}[y^*]}{\mathbb{E}[y]} \frac{\lim_{\rho \to 0} wn}{y^*} = H(0) \cdot 0 = 0. \tag{A.54}$$

For the limit at one, we note

$$\tilde{H}(\rho) \frac{\mathbb{E}[y]}{\mathbb{E}[y^*]} y^* = -\frac{w^2}{a} \left(\frac{1-\xi}{\xi} + \frac{1-\mu}{\mu} \frac{1}{\log \rho} \right) \overset{\rho \to 1}{\longrightarrow} +\infty \tag{A.55}$$

since $w \to a$ as $\rho \to 1$. Therefore, the optimal ρ when y/c is chosen optimally has the same properties that we used before holding y/c constant. In particular, the optimal replacement rate is unique and strictly below one. Also, we have $d\tilde{H}/d\rho > dI/d\rho$ at the optimal ρ as in Proposition (A.4). Taking the total derivative again, we have

$$\frac{d\rho}{da} = -\frac{\frac{\partial I}{\partial a} - \frac{\partial \tilde{H}}{\partial a}}{\frac{\partial I}{\partial \rho} - \frac{\partial \tilde{H}}{\partial \rho}} \tag{A.56}$$

where the denominator of the fraction is negative, so $d\rho/da$ has the same sign as its enumerator. The derivatives of I and $\tilde{H}$ with respect to productivity a are:

$$\begin{aligned} \frac{\partial I}{\partial a} &= \frac{\partial I}{\partial n} \frac{\partial n}{\partial a} \\ &= \frac{\partial n}{\partial a} \left(\frac{1-\rho}{n + (1-n)\rho} \right)^2 \left(\frac{w}{a} \frac{1-\mu}{\mu} \frac{1}{\log \rho} - \frac{1}{1-\rho} \right) < 0 \end{aligned} \tag{A.57}$$

and

$$\frac{\partial \tilde{H}}{\partial a} = \frac{\tilde{H}}{a} > 0. \tag{A.58}$$

Therefore $d\rho/da < 0$. ☐

BIBLIOGRAPHY

Adam, K., Marcet, A., 2011. Internal Rationality, Imperfect Market Knowledge and Asset Prices. Journal of Economic Theory, 146 (3), 1224–1252.

Adam, K., Beutel, J., Marcet, A., 2014. Stock Price Booms and Expected Capital Gains. Working Paper .

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

Adam, K., Marcet, A., Nicolini, J. P., 2013. Stock Market Volatility and Learning. Working Paper 336, Barcelona Graduate School of Economics.

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

Artus, P., Benassy-Quere, A., Boone, L., Cailloux, J., Delpla, J., Farhi, E., Gourinchas, P., Tirole, J., Wolff, G., 2013. Completing the Euro. Les notes du Conseil d'Analyse Economique.

Baily, M. N., 1978. Some Aspects of Optimal Unemployment Insurance. Journal of Public Economics, 10 (3), 379–402.

Balta, N., Delgado, J., 2009. Home Bias and Market Integration in the EU. CESifo Economic Studies, 55 (1), 110–144.

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

Barberis, N., Greenwood, R. M., Jin, L. J., Shleifer, A., 2015. X-Capm: An Extrapolative Capital Asset Pricing Model. Journal of Financial Economics, 115 (1), 1–24.

Bargain, O., Dolls, M., Fuest, C., Neumann, D., Peichl, A., Pestel, N., Siegloch, S., 2013. Fiscal Union in Europe? Redistributive and Stabilising Effects of an EU Tax-Benefit System. Economic Policy, 28 (75), 375–422.

Barro, R., 1990. The stock market and investment. Review of Financial Studies, 3 (1), 115–131.

Beaudry, P., Portier, F., 2006. Stock Prices, News, and Economic Fluctuations. American Economic Review, 96 (4), 1293–1307.

Beaudry, P., Portier, F., 2007. When Can Changes in Expectations Cause Business Cycle Fluctuations in Neo-classical Settings? Journal of Economic Theory, 135 (1), 458–477.

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

Bernanke, B. S., Gertler, M., Gilchrist, S., 1999. The financial accelerator in a quantitative business cycle framework. In J. B. Taylor, M. Woodford eds., Handbook of Macroeconomics, 1, Chapter 21, 1341–1393.

Bernoth, K., Engler, P., 2013. A Transfer Mechanism as a Stabilization Tool in the EMU. DIW Economic Bulletin, 3 (1), 3–8.

Blanchard, O. J., Jaumotte, F., Loungani, P., 2014. Labor market policies and IMF advice in advanced economies during the Great Recession. IZA Journal of Labor Policy, 3 (1), 1–23.

Blanchard, O., Rhee, C., Summers, L., 1993. The stock market, profit, and investment. Quarterly Journal of Economics, 108 (1), 115–136.

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

Bordo, M., Markiewicz, A., Jonung, L., 2011. A Fiscal Union for the Euro: Some Lessons from History. Working Paper 17380, NBER.

Brenke, K., 2013. A Skeptical View of Mechanisms for Business Cycle Harmonization in the Euro Area. DIW Economic Bulletin, 3 (1), 16–20.

Burda, M., Wyplosz, C., 1994. Gross worker and job flows in Europe. European Economic Review, 38 (6), 1287–1315.

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

Caputo, R., Medina, J. P., Soto, C., 2010. The Financial Accelerator Under Learning and The Role of Monetary Policy. Working Paper 590, Central Bank of Chile.

Cecchetti, S. G., Genberg, H., Wadhwani, S., 2002. Asset Prices in a Flexible Inflation Targeting Framework. Working Paper 8970, NBER.

Christoffel, K., Kuester, K., Linzert, T., 2009. The role of labor markets for euro area monetary policy. European Economic Review, 53 (8), 908–936.

Cochrane, J., 1992. Explaining the variance of price-dividend ratios. Review of Financial Studies, 5 (2), 243–280.

Cochrane, J. H., 2011. Presidential Address: Discount Rates. Journal of Finance, 66 (4), 1047–1108.

Coibion, O., Gorodnichenko, Y., 2010. Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts. Working Paper 16537, NBER.

Corbo, V., Osbat, C., 2013. Trade adjustment in the European Union - a structural estimation approach. Working Paper Series 1535, European Central Bank.

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

Covas, F., Den Haan, W. J., 2012. The Role of Debt and Equity Finance Over the Business Cycle. Economic Journal, 122, 1262–1286.

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

Elsby, M. W. L., Hobijn, B., Sahin, A., 2013. Unemployment Dynamics in the OECD. Review of Economics and Statistics, 95 (2), 530–548.

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

Eusepi, S., Preston, B., 2011. Expectations, Learning, and Business Cycle Fluctuations. American Economic Review, 101 (6), 2844–2872.

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

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

Farhi, E., Werning, I., 2012. Fiscal Unions. Working Paper 18280, NBER.

Filardo, A. J., 2001. Should Monetary Policy Respond to Asset Price Bubbles? Some Experimental Results. Research Working Paper RWP 01-04, Federal Reserve Bank of Kansas City.

Flynn, E., Crewson, P., 2009. Chapter 11 Filing Trends in History and Today. American Bankruptcy Institute Journal, 14, p. 65.

Fredriksson, P., Holmlund, B., 2006. Improving incentives in unemployment insurance: A review of recent research. Journal of Economic Surveys, 20 (3), 357–386.

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

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

Gourio, F., 2012. Disaster Risk and Business Cycles. American Economic Review, 102 (6), 2734–2766.

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

den Haan, W. J., Kaltenbrunner, G., 2009. Anticipated growth and business cycles in matching models. Journal of Monetary Economics, 56 (3), 309–327.

Hagedorn, M., Manovskii, I., 2008. The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited. American Economic Review, 98 (4), 1692–1706.

Hosios, A. J., 1990. On the Efficiency of Matching and Related Models of Search and Unemployment. Review of Economic Studies, 57 (2), 279–298.

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

Jaimovich, N., Rebelo, S., 2009. Can News about the Future Drive the Business Cycle? American Economic Review, 99 (4), 1097–1118.

Jung, P., Kuester, K., 2015. Optimal Labor-Market Policy in Recessions. American Economic Journal: Macroeconomics, 7 (2), 124–56.

Kiley, M., 2003. How Should Unemployment Benefits Respond to the Business Cycle? The B.E. Journal of Economic Analysis & Policy, 3 (1), 1–22.

Kimball, M. S., Fernald, J. G., Basu, S., 2006. Are Technology Improvements Contractionary? American Economic Review, 96 (5), 1418–1448.

Landais, C., Michaillat, P., Saez, E., 2015. A Macroeconomic Theory of Optimal Unemployment Insurance. Working Paper .

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

LEDUC, S., DEDOLA, L., CORSETTI, G., 2009. Should Monetary Policy React to Current Account Imbalances? Meeting Papers 1219, Society for Economic Dynamics.

LETTAU, M., UHLIG, H., 2000. Can Habit Formation Be Reconciled with Business Cycle Facts? Review of Economic Dynamics, 3 (1), 79–99.

LIU, Z., WANG, P., ZHA, T., 2013. Land-Price Dynamics and Macroeconomic Fluctuations. Econometrica, 81 (3), 1147–1184.

MAUSSNER, A., 2010. The Analytics of New Keynesian Phillips Curves. Discussion Paper Series 313, Universität Augsburg, Institute for Economics.

MEHRA, R., PRESCOTT, E. C., 1985. The equity premium: A puzzle. Journal of Monetary Economics, 15 (2), 145–161.

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

MITMAN, K., RABINOVICH, S., 2015. Optimal unemployment insurance in an equilibrium business-cycle model. Journal of Monetary Economics, 71, 99–118.

MOYEN, S., STÄHLER, N., 2014. Unemployment Insurance And The Business Cycle: Should Benefit Entitlement Duration React To The Cycle? Macroeconomic Dynamics, 18 (3), 497–525.

MUNDELL, R. A., 1961. A Theory of Optimum Currency Areas. American Economic Review, 51 (4), 657–665.

PERSSON, T., TABELLINI, G., 1996. Federal Fiscal Constitutions: Risk Sharing and Moral Hazard. Econometrica, 64 (3), 623–46.

PETRONGOLO, B., PISSARIDES, C. A., 2001. Looking into the Black Box: A Survey of the Matching Function. Journal of Economic Literature, 39 (2), 390–431.

VAN ROMPUY, H., 2012. Towards a Genuine Economic and Monetary Union. Report by the President of the European Council EUCO 120/12.

SANCHEZ, J. M., 2008. Optimal state-contingent unemployment insurance. Economics Letters, 98 (3), 348–357.

SARGENT, T. J., 2008. Rational Expectations. In D. R. Henderson ed., The Concise Encyclopedia of Economics, Library of Economics and Liberty, 2nd edition.

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

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

Shimer, R., 2005. The Cyclical Behavior of Equilibrium Unemployment and Vacancies. American Economic Review, 95 (1), 25–49.

Smets, F., Wouters, R., 2003. An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area. Journal of the European Economic Association, 1 (5), 1123–1175.

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

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

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