Slow Learning

Lawrence J. Christiano Martin Eichenbaum
Benjamin K. Johannsen*

February 27, 2026
Preliminary and Incomplete

Abstract

This paper analytically characterizes the speed of convergence under learning to a rational expectations equilibrium (REE) for a large class of multivariate models in which people's beliefs about model outcomes are central determinants of those outcomes. The paper then investigates what features of an economy determine the speed convergence. We do so by applying our analytic results to variants of the simple New Keynesian model when the Zero Lower Bound is and is not binding, and a medium scale DSGE model. Under certain circumstances, convergence of a learning equilibrium to the REE equilibrium can be so slow that analysis based on rational expectations is very misleading.

*The analysis and conclusions set forth are those of the authors and do not indicate concurrence by the Board of Governors or anyone else. The authors thank Klaus Adam, George-Marios Angeletos, Ed Herbst, Albert Marcet, Michael Massman, Alexander Mayer, Damjan Pfajfar, Bruce Preston, Tom Sargent, Robert Tetlow, Ivan Werning, Mike Woodford and participants in a number of seminars for helpful comments and suggestions.

1 Introduction

Rational expectations may be a useful modeling strategy in tranquil times like the Great Moderation. This strategy is less appealing when people are confronted with novel events, such as the Financial Crisis, the COVID-19 pandemic and the aftermath of the Russian invasion of the Ukraine. This paper analyzes the speed of learning and evolution of economic aggregates after a novel event. We assume that people must learn about their environment by forming beliefs about future economic outcomes and updating those beliefs as the data come in. Our analysis focuses on characterizing the speed of convergence of a learning equilibrium to a rational expectations equilibrium (REE). The critical issue is whether the speed of convergence is fast enough to render the REE a useful guide for normative and positive analyses after unusual events.

To address this question, we analytically characterize the speed of convergence in a broad class of non-stochastic learning models. When we confine ourselves to linear solutions of a model, our results apply to a broader class models, including medium-scale dynamic stochastic general equilibrium (DSGE) models, like CEE (2005). In all cases, the models that we consider have two important characteristics. First, people learn using either standard Bayesian methods or least-squares. Second, people's beliefs about model outcomes are central determinants of equilibrium outcomes.

We prove two propositions, which, taken together, show that a particular scalar parameter, b, of the multivariate, non-linear system determines the asymptotic speed of convergence to an REE.[1] That parameter, which also determines the E-stability of an REE ($b < 1$), can be calculated from the model solution. A model exhibits slow learning when b is less than but close to one. We investigate the economic determinants of that scalar variable, i.e., whether learning is fast or slow.

Our central finding is that when beliefs are partially self-fulfilling, learning equilibria converge slowly to rational expectations. Indeed, learning can be extraordinarily slow. For example, in the simple NK model, when wages and prices are sticky progress is measured in decades. A similar result holds for the benchmark CEE (2025) model. In the absence of investment adjustment costs, progress to the REE, meaningful convergence is measured in *centuries*. Even more dramatically, in the simple NK model when the zero lower bound (ZLB) is binding, the crucial parameter b is very close to 1, and progress is measured in *millennia*. In all three of these cases, positive and normative analyses based on rational expectations can be very misleading. As Vives (1993) writes, in a changing world, for all

[1] Our results are consistent with those in Christopeit and Massmann (2018), who study a univariate, linear, stochastic learning model.

practical purposes, "'slow' convergence may mean no convergence."

Consider the simple NK model in which both wages and prices are sticky. Suppose that at time zero people expect the mean of next period's consumption to be lower and the mean of inflation to be higher than in the REE. High expected inflation in future periods leads firms to set higher prices and households to set higher nominal wages at time zero. This response exerts upwards pressure on initial period inflation. Low expected aggregate consumption in the next period makes households feel poorer today because their income will be lower than in an REE.[2] This negative wealth effect induces households to cut back on consumption and supply more labor in the current period. So real wages fall in the current period, exerting downward pressure on inflation. Other things equal, time zero consumption is low. Because of competing forces, time zero inflation could be high or low relative to the REE value. Suppose that the impact of high expected inflation dominates the impact of the fall in real wages on current inflation. Then the monetary authority would set high nominal and real interest rates. As long as monetary policy gives positive weight to the output gap, there is a tradeoff between inflation and output objectives. The net result of that tradeoff is that at time zero, the real interest rate does not change by very much, so time 0 consumption is low and time 0 inflation is high relative to their values in an REE. So, the initial expectations of low consumption and high inflation are self confirming. However, they are not completely self-confirming because consumption is not quite as low and inflation is not quite as high as peoples' expected time period one values.

When agents update their expectations in the next period, they will do so based on the low equilibrium value of consumption and the high equilibrium value of inflation at time zero. So the expected value of future consumption and inflation will again be lower and higher, respectively, than their REE values. This mechanism repeats itself at $t = 2$ and in subsequent periods. The result is that expectations about consumption and inflation are partially self-fulfilling and learning is slow (b is high). Similar intuition obtains in the DSGE model of CEE (2005). The model of Smets and Wouters is the same as CEE (2005) up to the specification of shocks. It follows that learning will also be slow in that model. We conclude that learning is slow in estimated empirically plausible NK models.

The intuition for slow learning is particularly clear for the simple NK model when the ZLB is binding. Suppose that firms and households expect lower inflation in the future. Because of price-setting adjustment costs, firms are incentivized to cut prices today. In the ZLB, low inflation expectations mean households believe the real interest rate is high. Consequently, households reduce their demand for consumption, leading to a fall in the

[2] This effect is reinforced by households' expectations, that relative to the REE, they will receive less profits from monopolistic firms in the next period.

marginal cost of production. So, the actions of both households and firms lead to lower current inflation, consistent with their initial beliefs. With learning, low current inflation shifts expected inflation down in the next period. The previous mechanism repeats itself in the next period so that actual inflation in the next period is also low. We conclude that, in the ZLB, deflation expectations are partially self-fulfilling, and learning is slow (b is high).[3]

The speed of convergence plays a crucial role in analyzing the efficacy of various government policies like an increase in government purchases. These results emerge most starkly for the NK model when the ZLB is binding.[4] We use a non-linear version of the NK model in which only prices are sticky. This model has been widely used in the literature related to government purchases while the ZLB is binding. We also incorporate internalized learning for this analysis. In this model the effects of a rise in government purchases are much smaller under learning than under rational expectations. Under rational expectations, the government spending multiplier is very large when the ZLB binds because an increase in government purchases raises expected inflation (see Christiano et al. (2011)). Because the nominal interest rate is fixed, this rise generates a fall in the real interest rate, a rise in consumption, and a multiplier substantially larger than unity. Under learning, expected inflation is partially backward-looking and doesn't move much after an increase in government purchases. So, the real interest rate doesn't fall by very much, the key driver of the large REE multiplier is effectively eliminated, and the multiplier is close to unity.

We also consider the effects of monetary policy in the form of forward guidance when the ZLB is binding. To convey intuition as transparently as possible, we consider a simple form of forward guidance: the monetary authority commits to keeping the nominal interest rate at zero for one period after the shock that makes the ZLB binding returns to its steady-state level. The number of REEs proliferates under forward guidance, but only one REE is stable under learning. Consistent with the existing literature (for example, Del Negro et al. (2023) and Woodford (2012)), we find that even this simple form of forward guidance is powerful under rational expectations. As is well known, the power of forward guidance under rational expectations reflects its strong effect on expected inflation. Under learning, expectations are partially backward-looking, and forward guidance is not very powerful. So, as with fiscal policy, a REE-based analysis of monetary policy can be very misleading.

We show that our key conclusions are robust to a variety of perturbations. First, as noted above, those conclusions hold in empirically plausible DSGE models like the model of

[3]These results are consistent with results in Heemeijer et al. (2009) and Hommes (2011), who analyzed the interactions between beliefs and outcomes in an experimental setting.

[4]Much of the work in the initial aftermath of that event combined rational expectations with the NK model. See, for example, Eggertsson and Woodford (2004), Christiano et al. (2011) and Del Negro et al. (2023).

CEE (2005). Second, our results are robust to whether agents solve their problems using a version of Kreps (1998)'s *Anticipated Utility* (our benchmark analysis) or *internally rational* learning. Under anticipated utility, people update their beliefs every period as new data comes in. But, when they make their decisions, people proceed as though their beliefs will never be revised again.[5] Under internally rational learning, people fully integrate the fact that they are learning when they solve their problems, i.e., households and firms are internally rational in the sense defined by Adam and Marcet (2011). Implementing internal rationality in the non-linear solution of the model is computationally very challenging.[6] The associated computational burden explains why much of the learning literature works with a version of Kreps (1998)'s *Anticipated Utility* approach.[7] Because of computational challenges, we investigate the implications of the two different learning models using the full non-linear solution to the simple NK model when the ZLB is binding. Qualitatively similar results under Krep's anticipated utility and internal rational learning.

Third, we show that our results are robust to whether we work with the fully non-linear or linear approximations to model solutions. Again, for computational reasons it is difficult to work with a fully nonlinear version of the DSGE model in CEE (2005). So we use the simple NK model when the ZLB is binding. In that model the speed of convergence to the REE is also particularly slow.

Fourth, we investigate whether our results regarding the asymptotic rate of convergence to the REE are useful for characterizing convergence speed over short horizons. Based on our analysis of the NK model, we show that the answer is yes.

2 Related Literature (Incomplete)

Our paper is related to several literatures. The first is a literature that studies the conditions under which non-stochastic learning equilibria converge to an REE (see Evans and Honkapohja (2000)). In contrast, we study the rate of convergence to an REE in that class of models.

The second is the literature that studies the properties of recursive stochastic estimators in learning models. Ljung (1977) establishes that a recursive estimator, $\widehat{\theta}_t$, converges almost surely to a limiting value, θ, if a particular ordinary differential equation (ODE), determined

[5] This approach has been criticized for its internal inconsistency (see Cogley and Sargent (2008) and Adam and Marcet (2011)).

[6] For example, we solve the simple NK model when the ZLB is binding using a compiled programming language (c++) and we make use of more than 300 processors. See Appendix C for details.

[7] This approach has been criticized for its internal inconsistency (see Cogley and Sargent (2008) and Adam and Marcet (2011)).

by the economic model, has eigenvalues with real parts that are less than unity. Marcet and Sargent (1989b; 1989a), Woodford (1990), Evans and Honkapohja (2000; 2001), and others build on Ljung (1977) to study the conditions under which learning equilibria converge to an REE. Marcet and Sargent (1995) numerically study the rate at which these learning equilibria converge to an REE. Christopeit and Massmann (2018) provide analytic results in a linear, scalar stochastic model. In contrast, we provide multivariate, non-stochastic results and apply them to the non-linear NK model.

Ferrero (2007) discusses learning in the context of a linear NK model in which the ZLB on interest rates is not binding. He uses the simulation methods proposed by Marcet and Sargent (1995) to study convergence rates of learning equilibria. Ferrero (2007) adopts the so-called Euler-equation approach to learning as opposed to our approach; see Evans (2021) for a definition of the Euler-equation approach to learning, and see Preston (2005) and Adam and Marcet (2011) for a critique of that approach.[8] Another difference with Ferrero (2007) is that we compare convergence rates in a non-linear NK model when the ZLB on interest rates is and is not binding.

Cogley and Sargent (2008), Adam and Marcet (2011), and Adam et al. (2017) numerically analyze endowment economies in which people learn and make decisions in an internally rational way. Adam and Merkel (2019) use this approach to numerically analyze a real business cycle model in which peoples' beliefs do not nest an REE. In contrast, we provide analytic results about rates of convergence for a broad class of models and study a non-linear NK model in which peoples' beliefs do nest an REE.

Preston (2005) and Eusepi et al. (2022) use the anticipated utility approach to study the effects of monetary policies in linearized NK models under learning. In contrast, we work with a non-linear model in which people make internally rational decisions. ding fundamentals that are observed with noise.

A different literature investigates the information content in prices regarding fundamentals that are observed with noise. In that context, Vives (1993) asks: how quickly do people's beliefs about an exogenous cost parameter converge? A large literature also explores the speed with which people learn the parameters of *exogenous* stochastic processes. For example, Erceg and Levin (2003), Gust et al. (2018), and Farmer et al. (2021) describe an empirically relevant set of time series representations with hard-to-learn low-frequency components. In contrast, we study convergence rates for beliefs about objects whose values depend on those beliefs.

Heemeijer et al. (2009) and Hommes (2011) study positive and negative feedback loops from expectations to outcomes using laboratory experiments and univariate models with

[8]Our approach is an example of what Evans (2021) calls the agent-based approach to learning.

constant gain. In contrast, we analytically characterize convergence rates of beliefs for a broad class of models and numerically analyze rates of convergence in a multivariate NK model under Bayesian learning.

Our paper is also related to a recent game-theoretic grounded literature that analyzes the implications of bounded rationality for the effectiveness of fiscal and monetary policy. Farhi and Werning (2019) use k-level thinking models to study how deviations from rational expectations affect the effectiveness of forward guidance. García-Schmidt and Woodford (2019) study forward guidance and interest rate pegs using reflective expectations. Iovino and Sergeyev (2023) apply k-level thinking and reflective expectations to analyze the effects of quantitative easing. Angeletos and Lian (2017) develop the idea that a lack of common knowledge can attenuate general-equilibrium effects and damp the effects of government spending. Angeletos and Lian (2017; 2018) analyze the consequences of bounded rationality for the size of fiscal multipliers.

Farhi and Werning (2019), Farhi et al. (2020) and Woodford and Xie (2019; 2022) use different models of bounded rationality to study the size of the government-spending multiplier. Vimercati et al. (2021) assess the implications of bounded rationality for the effectiveness of tax and government spending policy at the ZLB. They do so through the lens of a standard NK model in which people are dynamic k-level thinkers.

In all of the papers just cited, individuals have a limited ability to understand the general equilibrium consequences of monetary and fiscal policies. Like learning, this type of deviation from rational expectations can limit the power of forward guidance. Our paper studies a form of deviation from rational expectations different from those cited in the previous two paragraphs. Moreover, in contrast to our analysis, these papers do not analyze rates of convergence to rational expectations.

3 Learning in a Non-Linear Environment

Here, we consider the speed of convergence of learning in the following non-linear environment, which is very similar to that studied in Evans and Honkapohja (2000) who establish conditions for convergence of beliefs to REE. In contrast, we focus on the rate of convergence. Where possible, we use the notation of Evans and Honkapohja (2000).

Let θ_t be a $k-$dimensional vector of variables which summarizes people's period t priors about a set of variables that will be determined at time $t+1$. We interpret θ_t as a deviation from a particular fixed point of beliefs in our learning algorithm.[9] Given a set of beliefs, θ_{t-1}, the environment generates outcomes in period t according to the non-linear function,

[9] We do not require that this fixed point be unique.

$M(\theta_{t-1}, \gamma_t)$. The vector, θ_t, evolves according to

$$\theta_t = \theta_{t-1} + \gamma_t \left[M(\theta_{t-1}, \gamma_t) - \theta_{t-1}\right] \tag{1}$$

for $t = 1, 2, 3, \ldots$. Here, θ_0 is given and $\gamma_t = \frac{1}{c_1+t}$ for $c_1 \geq 0$ is the gain. This type of gain parameter emerges from standard Bayesian learning as well as least squares learning. Following Evans and Honkapohja (2000), we assume that the vector-valued function $M : \mathbb{R}^k \times \mathbb{R} \to \mathbb{R}^k$ has the following properties: $M(0,0) = 0$; M is continuously differentiable in a neighborhood of the origin; $M(0, \gamma_t) = 0$; and $M(0, \gamma_t)$ is continuously differentiable in a neighborhood of $(0, \gamma_t)$. Let D_1M denote the derivative of M with respect to the vector θ_{t-1} (which means it is a $k \times k$ matrix). We assume the real parts of the eigenvalues of D_1M are strictly less than unity. The scalar, b, denotes the largest real part of the eigenvalues of D_1M.

We consider the same environment as Evans and Honkapohja (2000), but with a more restrictive specification of γ_t.[10] They establish that for, (i) $b < 1$, and (ii) θ_0 sufficiently close to 0, θ_t converges to zero. We follow their technique to study the rate at which θ_t converges to zero. Christopeit and Massmann (2018) consider a scalar, linear, stochastic version of our environment. Their results imply that for the non-stochastic environment studied here, θ_t asymptotically converges to zeto at the rate t^{b-1}. We extend this result to a multivariate, non-linear environment.[11]

To analyze rates of convergence it is convenient to adopt the following definition.

Definition 1. For $b < 1$, we say that $x_t \simeq t^{b-1}$ if for any $0 < \delta$, (1) $\lim_{t\to\infty} \frac{\|x_t\|}{t^{b-1+\delta}} = 0$, and (2) $\lim_{t\to\infty} \frac{\|x_t\|}{t^{b-1-\delta}} = \infty$. If $x_t \simeq t^{b-1}$ then we say that x_t *asymptotically converges to zero at the rate* t^{b-1}.

Here, $\|\cdot\|$ denotes a norm on $\mathbb{R}^k$. The first part of definition 1 says that for any positive δ, $\|x_t\|$ asymptotically converges no slower than $t^{b-1+\delta}$. The second part says that $\|x_t\|$ asymptotically converges no faster than $t^{b-1-\delta}$. In this sense, b characterizes the power rate of convergence. Importantly, two series that asymptotically converge to zero at the rate t^{b-1} may behave very differently for finite T. Some examples of series that are different even for large t, even though they asymptotically converge to zero at the rate t^{b-1}, include

[10] They assume $1 > \gamma_t > 0$, $\gamma_t \to 0$ and $\lim_{T\to\infty} \sum_{t=1}^{T} \gamma_t$ diverges. Our specification of γ_t satisfies these conditions.

[11] Ljung (1977) showed that to determine whether θ_t converges, it is useful to consider the ordinary differential equation, $\dot{\theta}(\tau) = D_1M\theta(\tau) - \theta(\tau)$, where τ evolves in continuous time. In the scalar case, $b = D_1M$ and the solution to this equation is $\theta(\tau) = \left[e^{(b-1)}\right]^{\tau} \theta(0)$. The same parameter, b, determines whether $\theta(\tau)$ and θ_t converge as well as their rates of convergence. But, the rates of convergence of these variables are qualitatively different: $\theta(\tau)$ converges at a geometric rate in $\tau-$time while θ_t converges at a power rate in $t-$time. See the appendix for further discussion in the context of a linear model.

$x_t = \log(t)\, t^{b-1}$, $y_t = t^{b-1}/\log(t)$, and $z_t = [2 + \sin(t)]\, t^{b-1}$. These series asymptotically converge at the same rate when considering only rates of power convergence, which is what our definition captures. There are, of course, series that converge at a faster rate than power convergence, such as those that converge at a geometric rate (for example, $x_t = \rho^t$ for $|\rho| < 1$). As it turns out, the NK model with Bayesian learning (or any model whose reduced form satisfies equation 1) rate exhibits power convergence.

We now state two propositions that are proved in the appendix. The first establishes that there exists a neighborhood of the origin denoted by $U \in \mathbb{R}^k$ so that for every $\theta_0 \in U$ the implied sequence, θ_t, converges to the origin and satisfies part (1) of definition 1. The second proposition establishes that there exists a $\theta_0 \in U$ that implies a sequence, θ_t, that satisfies part (2) of definition 1. Both propositions are stated under the assumptions related to equation (1) that are stated above and under the assumption that $b < 1$.

Proposition 1. *There exists a neighborhood U of 0 such that for any $0 < \delta$ if $\theta_0 \in U$ then $\lim_{t\to\infty} \|\theta_t\| = 0$ and $\lim_{t\to\infty} \frac{\|\theta_t\|}{t^{b-1+\delta}} = 0$.*

Proposition 2. *For any $0 < \delta$ and any neighborhood U of 0, there exists a $\theta_0 \in U$ so that $\lim_{t\to\infty} \frac{\|\theta_t\|}{t^{b-1-\delta}} = \infty$.*

The previous propositions establish that there exists a θ_0 near the origin that generates a sequence, $\theta_t \simeq t^{b-1}$. They also imply that there is a neighborhood of the origin in which there is no θ_0 that generates a sequence, θ_t, that converges to zero at a rate slower than t^{b-1}.

In the remainder of the paper, we apply and discuss the implications of these propositions for the New Keynesian model.

4 The speed of learning in the simple NK model

In this section we study the rate of convergence of beliefs in the NK model. We establish that, when the ZLB is not binding, if nominal prices are sticky but nominal wages are perfectly flexible, then convergence to the REE is fast. However if nominal wages are sticky, the rate of convergence is very slow regardless of whether nominal prices are sticky or not. Moreover, there is an important interaction between sticky wages and sticky prices in slowing down the rate of convergence to the REE

4.1 The simple NK model with sticky nominal wages and prices

In this subsection we describe a version of the simple NK model in which both nominal wages and prices are sticky. We model nominal rigidities as arising from Rotemberg-style adjustment costs.

4.1.1 Firms' Problems

A final homogeneous good, Y_t, is produced by competitive and identical firms using the technology

$$Y_t = \left(\int_0^1 Y_{f,t}^{\frac{\varepsilon-1}{\varepsilon}} df \right)^{\frac{\varepsilon}{\varepsilon-1}}, \tag{2}$$

where $\varepsilon > 1$. The representative homogeneous-good firm chooses inputs, $Y_{f,t}$, to maximize profits $Y_t P_t - \int_0^1 Y_{f,t} P_{f,t} df$, subject to equation 2. The representative firm's first-order condition for the f^{th} input is

$$Y_{f,t} = \left(\frac{P_{f,t}}{P_t} \right)^{-\varepsilon} Y_t. \tag{3}$$

At the beginning of a period, before markets open, monopolistic firms also know the value of the vector, Θ_t, which summarizes their beliefs about the distribution of the vector, x_t

$$x_t = \begin{bmatrix} C_t, & \pi_t, & w_t, & w_{t-1} \end{bmatrix}'. \tag{4}$$

The variables C_t, π_t, and w_t denote aggregate consumption, aggregate inflation, and the economy-wide real wage rate at time t. The variable, π_t, corresponds to P_t/P_{t-1}, where P_t and P_{t-1} are the time t and $t-1$ values of aggregate price level. We include w_{t-1} in x_t because w_{t-1} is a state variable in the REE when there are nominal wage rigidities. The vector Θ_t denotes a set of variables which summarizes the firm's period t priors about the future values of a set of aggregate variables. In principal, firms could internalize the effect of x_t on their beliefs about the distribution x_{t+1}. Those beliefs, Θ_{t+1}, are given by

$$\Theta_{t+1} = L\left(\Theta_t, x_t\right). \tag{5}$$

The form of L depends on the assumed model of learning and decision making.

In a standard recursive equilibrium, people know current-period market prices and profits when they make their current decisions. Typically, when markets open in these models, people can deduce the values of all aggregate variables prices and profits from a small set of variables. In our context this set of variables is given by x_t. We assume that firms observe x_t when they make their time t price decision.

The f^{th} intermediate good is produced by a monopolist with production technology $Y_{f,t} = N_{f,t}$, where $N_{f,t}$ is labor hired by firm f. Let $p_{f,t} = P_{f,t}/P_t$. Also, let p_f' denote the firm's current choice of price scaled by the current aggregate price index. In our scaled notation,

$$\frac{p_f'}{p_f}\pi = \frac{P_{f,t}}{P_{f,t-1}}. \tag{6}$$

Firms value a unit of real profits by the representative household's marginal utility of consumption, $1/C_t$. Nominal prices are sticky as in Rotemberg (1982). The real price set by firm f , $p_{f,t} = p_f\left(p_{f,t-1}, \Theta_t, x_t\right)$, satisfies

$$\begin{aligned} p_f\left(p_{f,t-1}, \Theta_t, x_t\right) = \text{argmax}_{p_{f,t}} \frac{1}{C_t}\left\{\left(p_{f,t} - \left(1-\tau_f\right) w_t\right) p_{f,t}^{-\varepsilon} Y_t - \frac{\phi}{2}\left(\frac{p_{f,t}}{p_{f,t-1}}\pi_t - 1\right)^2 C_t\right\} \\ + \beta \mathbb{E}_t V_f\left(p_{f,t}, \Theta_t, x_{t+1}\right). \end{aligned} \tag{7}$$

Because firms see P_t, choosing $p_{f,t}$ is equivalent to choosing $P_{f,t}$. Note that the adjustment costs pertain to nominal prices. In equation (7), we follow the literature by scaling price adjustment costs by real GDP.[12] Also, τ_f is a government tax subsidy on employment designed to eliminate the effect of monopoly distortions in steady state.[13] We assume that each firm knows the model's static equilibrium conditions so they can deduce Y_t from x_t. The expectation operator, $\mathbb{E}_t$, is evaluated using the marginal data density for x_{t+1} implied by Θ_t.

Here, $V_f\left(p_{f,t}, \Theta_{t+1}, x_{t+1}\right)$ appearing in (7) denotes the value of the firm's problem. The function, $V_f\left(p_{f,t-1}, \Theta_t, x_t\right)$ has the fixed-point property

$$\begin{aligned} V_f\left(p_{f,t-1}, \Theta_t, x_t\right) = \max_{p_{f,t}} \frac{1}{C_t}\left\{\left(p_{f,t} - \left(1-\tau_f\right) w_t\right) p_{f,t}^{-\varepsilon} Y_t - \frac{\phi}{2}\left(\frac{p_{f,t}}{p_{f,t-1}}\pi_t - 1\right)^2 C_t\right\} \\ + \beta \mathbb{E}_t V_f\left(p_{f,t}, \Theta_t, x_{t+1}\right). \end{aligned} \tag{8}$$

After choosing $p_{f,t}$, the firm supplies all of its good that is demanded at that price.

4.1.2 Anticipated utility

Virtually all of the related literature works with a version of the Kreps' *Anticipated Utility* approach to how people integrate learning into their decisions. In this approach, people update their beliefs every period as new data come in. But, when they make decisions, people proceed as though their beliefs will never be revised again. This approach has very significant computational advantages, so we use the anticipated utility approach in this subsection.

[12] See, for example, Kaplan and Violante (2018, page 711).

[13] That is, $\left(1-\tau_f\right)\varepsilon/\left(\varepsilon-1\right) = 1$.

This approach has been criticized for being internally inconsistent (see Cogley and Sargent (2008) and Adam and Marcet (2011)). In section 5.4, we assess the robustness of our results to using internalized learning to model firms' behavior. Under this approach, when making decisions, firms take into account uncertainty about the distribution of x and the fact that beliefs about that distribution will evolve as new data arrive (see section 4.1.2). So, firms are internally rational in the sense of Adam and Marcet (2011).

To implement the anticipated utility approach, we assume that firms believe $\Theta_{t+j} = \Theta_t$ for $j \geq 1$ when making their decisions state contingent decisions. They also ignore the fact that, after they see current x, they will update their views, using $\Theta_{t+1} = L(\Theta_t, x_t)$.

4.1.3 The Household's Problem

There is a continuum of households that has unit measure. We index the households by h. Let $C_{h,t}$, $b_{h,t}$, and $w_{h,t}$ denote household h's consumption, end-of-period bond holdings, and the real wage that the household set.

The household enters a period with a stock of bonds, $b_{h,t-1} = B_{h,t-1}/P_{t-1}$. Here, $B_{h,t-1}$ denotes the beginning-of-period t payoff on nominal bonds acquired in the previous period, when the price of consumption goods was P_{t-1}. At the beginning of a period, before markets open, the household also knows the value the vector, Θ, which summarizes its beliefs about the distribution of the vector, x.

We incorporate sticky nominal wages using the following variant of Erceg et al. (2000). We assume that a final unit of the labor input into monopolists' production function, N_t, is produced by a competitive firm using the technology

$$N_t = \left(\int_0^1 N_{h,t}^{\frac{\epsilon-1}{\epsilon}} dh \right)^{\frac{\epsilon}{\epsilon-1}}. \tag{9}$$

So, household h faces a demand curve for its labor of the form

$$N_{h,t} = \left(\frac{w_{h,t}}{w_t} \right)^{-\epsilon} N_t, \tag{10}$$

where w_t is the real price of N_t. The household solves the problem:

$$\max_{C_{h,t}, w_{h,t}, b_{h,t}} \log(C_{h,t}) - \frac{\chi}{2} N_{h,t}^2 - \frac{\phi_W}{2} \left(\frac{w_{h,t}}{w_{h,t-1}} \pi_t - 1 \right)^2 + \beta \mathbb{E}_t V_h(b_{h,t}, w_{h,t}, \Theta_t, x_{t+1}). \tag{11}$$

The scalar $\phi_W > 0$ controls the cost of adjusting nominal wages. Because households see

P_t, choosing $w_{h,t}$ is equivalent to choosing $W_{h,t}$. The household solves its problem subject to equation (10) and the budget constraint:

$$C_{h,t} + \frac{b_{h,t}}{R_t} + \frac{\vartheta}{2} b_{h,t}^2 \leq \frac{b_{h,t-1}}{\pi_t} + (1+\tau_w) w_{h,t} N_{h,t} + T_t. \tag{12}$$

Here, T_t denotes profits net of lump-sum taxes, R_t denotes the nominal rate of interest, and τ_w is a government tax subsidy on employment designed to eliminate the effect of monopoly distortions in steady state. The term $\frac{\vartheta}{2} b_{h,t}^2$ does not usually appear in the household budget constraint in the NK model. Including the term $\frac{\vartheta}{2} b_{h,t}^2$ has no effect on an REE because households know that $b_{h,t} = 0$. In a learning equilibrium, households will contemplate holding non-zero values of bonds. The term $\frac{\vartheta}{2} b_{h,t}^2$ can be thought of as a cost associated with holding non-zero values of bonds. The reason for including that term is analagous to the reason for including a similar costs to domestic households for holding foreign bonds in small-open-economy models (see Schmitt-Grohé and Uribe (2003)). In practice, we set ϑ to be a very small number and no costs are actually paid (even in learning equilibria) because $b_{h,t}$ has to be zero for bond markets to clear.

In equation (11), $V_h(b_{h,t}, w_{h,t}, \Theta_t, x_{t+1})$ denotes household $h's$ value function. The function, $V_h(b_{h,t-1}, w_{h,t-1}, \Theta_t, x_t)$, satisfies the following fixed point:

$$\begin{aligned} V_h(b_{h,t-1}, w_{h,t-1}, \Theta_t, x_t) = \max_{C_{h,t}, w_{h,t}, b_{h,t}} \Bigg\{ & \log(C_{h,t}) - \frac{\chi}{2} N_{h,t}^2 - \frac{\phi_W}{2} \left(\frac{w_{h,t}}{w_{h,t-1}} \pi_t - 1 \right)^2 \\ & + \beta \mathbb{E}_t V_h(b_{h,t}, w_{h,t}, \Theta_t, x_{t+1}) \Bigg\}, \end{aligned} \tag{13}$$

subject to equations (10) and (12). After choosing $w_{h,t}$, the household supplies all labor that is demanded at that wage.

Households and firms have the same information sets and update priors in the same way. So, the expectations operator, $\mathbb{E}_t$, in the household's problem is the same as the one in the firm's problem.

4.1.4 Monetary and Fiscal Policy

Monetary policy sets the gross nominal interest rate, R_t, according to

$$\log(R_t/R) = \alpha_\pi (\pi_t - 1) + \alpha_Y \log(Y_t/Y). \tag{14}$$

The government finances subsidies to households and firms with lump-sum taxes and balances its budget in each period.

4.1.5 Resource constraint and aggregate mappings

For individual households' and firms' problems to be well defined, they must know the values of seven aggregate variables, C_t,π_t, R_t, Y_t, N_t, w_t, T_t. We assume that each agent knows the model's static equilibrium conditions so they can deduce those variables from $x_t = \begin{bmatrix} C_t, & \pi_t, & w_t \end{bmatrix}'$. We denote this mapping by $F(x_t)$. Households and firms derive Y_t from x_t using the resource constraint

$$Y_t = C_t \left(1 + \frac{\phi}{2} (\pi_t - 1)^2 \right),$$

and derive R_t from x_t using equation 14. The mappings from x_t to Y_t, N_t, and w_t are given by

$$Y_t = C_t \left(1 + \frac{\phi}{2} (\pi_t - 1)^2 \right), \; N_t = Y_t.$$

These equalities correspond to goods market clearing and the aggregate production function. In the case that only wages are sticky, households and firms know that $w_t = 1$ in every period. In the case that only prices are sticky, households and firms know that $w_t = \chi C_t^{\sigma} N_t^{\zeta}$. These equalities hold in every period of our learning equilibria (described in the next sub-section). Aggregate firm profits net of taxes implied by x and r are

$$T_t = (1 - (1 - \tau_f) w_t) Y_t - \frac{\phi}{2} (\pi_t - 1)^2 C_t - (\tau_f + \tau_w) w_t Y_t.$$

In the case that wages are not sticky, $\tau_w = 0$.

4.1.6 Equilibrium Definitions

We now define a period equilibrium.

Definition 2. Given Θ_t, a period equilibrium is a set of values of x_t such that

(i) households and firms solve their optimization problems, defined in equations 7 and 13, respectively

(ii) labor, goods and bond markets clear

(iii) $p_{f,t} = 1$, $C_{h,t} = C_t$, $N_{h,t} = N_t$, $b_{h,t} = 0$, $w_{h,t} = w_t$

Condition (ii) reflects our assumption that households and firms are identical across f and h. Because firms are identical, in a learning equilibrium, no firm will ever inherit a $p_{f,t} \neq 1$.

Then, equation 6 and the first part of condition (iii) imply that people's views about inflation, π_t, are correct. The second and third parts of condition (iii) imply that people's views about C_t and N_t are correct. The condition that $p_{f,t} = 1$ imposes that that firms find it optimal to set $P_{f,t} = P_t$. The condition that $w_{h,t} = w_t$ imposes that households find it optimal to set their household-specific wage equal to the aggregate wage. The condition that $b_{h,t} = 0$ imposes that households all choose the same amount of bonds, which are in net zero supply.

We now define a learning equilibrium.

Definition 3. A *learning equilibrium* is a sequence of period equilibria in which beliefs are updated according to equation (5).

In a learning equilibrium, the value of Θ_t in the first period is exogenous. We assume that people's priors about the economic variables, x_t, are very diffuse. Below, we describe how our parameterization of the initial Θ_t captures this property.

4.2 Solving the model and specification of beliefs

We analyze the model's properties using a log-linear approximation to the non-linear learning equilibrium. In section 5.3, we assess the accuracy of log-linear approximations to the non-linear solution.

To specify beliefs about aggregate variables, we assume that people believe that

$$\begin{bmatrix} \widehat{C}_t \\ \widehat{\pi}_t \\ \widehat{w}_t \end{bmatrix} = \begin{bmatrix} \mu_C \\ \mu_\pi \\ \mu_w \end{bmatrix} + \begin{bmatrix} \omega_C \\ \omega_\pi \\ \omega_w \end{bmatrix} \widehat{w}_{t-1} + \begin{bmatrix} \varepsilon_{C,t} \\ \varepsilon_{\pi,t} \\ \varepsilon_{w,t} \end{bmatrix}. \tag{15}$$

Here, $\widehat{x}_t$ is the log deviation of the variable x_t from its value in the target-inflation REE steady state.[14] Also, $\varepsilon_{C,t} \sim N(0, \sigma_C^2)$, $\varepsilon_{\pi,t} \sim N(0, \sigma_\pi^2)$, and $\varepsilon_{w,t} \sim N(0, \sigma_w^2)$. These distributions are independent from each other and across time. People are uncertain about the values of μ_i, σ_i^2 for $i \in \{C, \pi, w\}$. To simplify the analysis, we assume that people know the REE values of ω_i and learning is only about μ_i. Presumably learning would be no faster if people also had to learn about ω_i. In this sense, our assumption that people do not learn about ω_i is conservative.

People's priors about μ_i conditional on σ_i^2 are Normal, parameterized with a mean, $m_{i,t}$, and variance, $\sigma_i^2/\lambda_{i,t}$, where $\lambda_{i,t}$ characterizes the precision of the prior about μ_i. The marginal density of their prior for σ_i^2 is proportional to an inverse-gamma distribution,

[14] In the appendix, we analyze the robustness of our results to linearizing around the deflation REE steady state. For the cases we consider, we find that learning is still slow.

with shape and scale parameters, $\alpha_{i,t}$ and $\left(\psi_{i,t}^2\left(\alpha_{i,t}+1/2\right)\right)$, respectively. The prior for σ_i^2 is not exactly inverse-gamma distributions because we truncate the support of σ_i^2 so that $\mathbb{E}_t\left[C_{t+1}\right]$, $\mathbb{E}_t\left[\pi_{t+1}\right]$, and $\mathbb{E}_t\left[w_{t+1}\right]$ have finite values. We find it convenient to express the scale parameter in this way because $\psi_{i,t}$ is a consistent estimator for σ_i. The joint density of μ_i, σ_i^2 is proportional to the Normal inverse-gamma distribution. The posterior distribution is also proportional to the Normal inverse-gamma distribution,. The function, L, in equation 5 can be constructed using standard updating formulas, which are detailed in appendix C. We assume that in the current and all future periods, firms believe that $\log(x)$ will be drawn from a Normal distribution with the mean and variance of the i^{th} variable in x_t fixed at the values of $m_{i,t}$ and $\psi_{i,t}^2$ from Θ_t. This assumption implies that firms, ignore uncertainty about the mean and variance of the distribution of $\log(x)$.

Because we work with the log-linear approximation, $\psi_{i,t}$ does not affect people's decisions. We set $\lambda_{i,0}=1$. Because $\lambda_{i,t}$ does not enter the linear approximation to to the solution of the model and only affects γ_t, we are able to summarize beliefs by

$$m_t = \begin{bmatrix} m_{C,t}, & m_{\pi,t}, & m_{w,t} \end{bmatrix}'. \tag{16}$$

When relevant, we discuss the exact values for m_0 below.

4.2.1 Connecting to Propositions 1 and 2

Because of symmetry among households and firms and the requirement that markets clear, in a learning equilibrium $\hat{w}_{h,t-1}=\hat{w}_{t-1}$, $\hat{b}_{h,t-1}=0$, and $\hat{p}_{f,t-1}=0$. Given these states and their beliefs, m_{t-1}, households and firms update their beliefs according to Bayes' rule so that

$$m_t = m_{t-1}+\gamma_t\left(\begin{bmatrix}\hat{C}_t\\ \hat{\pi}_t\\ \hat{w}_t\end{bmatrix}-\begin{bmatrix}\omega_C\\ \omega_\pi\\ \omega_w\end{bmatrix}\hat{w}_{t-1}-m_{t-1}\right)=m_{t-1}+\gamma_t\left(\mu(m_{t-1})-m_{t-1}\right). \tag{17}$$

The value of $\mu(m_{t-1})$ is computed using equation 15 by finding the three values $\varepsilon_{C,t}$, $\varepsilon_{\pi,t}$, and $\varepsilon_{w,t}$ that lead households and firms to make decisions so that markets clear, i.e. the three equations $\hat{b}_{h,t}=\hat{p}_{f,t}=\hat{w}_{h,t}-\hat{w}_t=0$ hold. Under our assumptions, households and firms know the values of $\hat{C}_t$, $\hat{\pi}_t$, and $\hat{w}_t$ when they make their decisions. Those decisions must collectively give rise to those value of $\hat{C}_t$, $\hat{\pi}_t$, and $\hat{w}_t$. Because household and firms don't have RE, they believe that the values of $\hat{C}_t$, $\hat{\pi}_t$, and $\hat{w}_t$ are determined by the shocks $\varepsilon_{C,t}$, $\varepsilon_{\pi,t}$, and $\varepsilon_{w,t}$, which they take to be exogenous. However, in the learning equilibrum, values of $\varepsilon_{C,t}$, $\varepsilon_{\pi,t}$, and $\varepsilon_{w,t}$ are endogenous and determined so that household and firm choices are consistent with the aggregate values for $\hat{C}_t$, $\hat{\pi}_t$, and $\hat{w}_t$.

The value of $\hat{w}_{t-1}$ does not affect the value of μ. To see why, note that intercept terms in a dynamic linear model only affect the intercept terms in the solution to the model (see Sims (2001)). Conditional on beliefs, we solve the linear approximation to the household and firm problems.[15] In an REE, the laws of motion for $\hat{b}_{h,t}$, $\hat{p}_{f,t}$, and $\hat{w}_{h,t} - \hat{w}_t$ have a coefficient of zero on $\hat{w}_{t-1}$. So, under our assumptions they also have a coefficient of zero on $\hat{w}_{t-1}$ in the learning equilibrium. As a result, the values of $\varepsilon_{i,t}$ that determine $\mu(m_{t-1})$ are not a function of $\hat{w}_{t-1}$.

Equation 17 is of the same form as equation 1 with $M(\theta_{t-1}, \gamma_t) = \mu(m_{t-1})$. So, propositions 1 and 2 apply to this system. We numerically calculate the maximal real part of the eigenvalues of $D\mu(0)$, which corresponds to b in propositions 1 and 2.

4.3 Speed of convergence

In the following subsection we analyze the speed of convergence using a calibrated version of the simple NK model. We assume the following benchmark parameter values:

$$\beta = 0.995, \varepsilon = 4, \epsilon = 4, \phi_P = 110, \phi_W = 110, \chi = 1, \alpha_\pi = 1.5, \alpha_Y = 0.25.$$

In the benchmark model, prices are flexible and only wages are sticky. We choose β so that the steady state real interest rate is 2 percent at an annual rate. We set $\epsilon = \varepsilon = 4$ implying a steady state markup of 33 percent in product and labor markets. These values are within the range of values considered in the related literature. It is well known that the linear approximation to the solution of simple NK model with Rotemberg-style sticky prices is observationally equivalent to that of the simple NK model with Calvo-style sticky prices. We choose the value of ϕ_P so that in the observationally equivalent Calvo model the probability of not changing price in any given quarter is 0.85. For symmetry, we set $\phi_W = \phi_P$. The value of χ is chosen so that in steady state labor supply is one. The values of α_π and α_Y are chosen to be standard values (see Taylor (1999)).

Table 1 reports the value of b in the benchmark model and in the following perturbations of that model: (i) wages are flexible and only prices are sticky ($\phi_P = 110$ and $\phi_W = 0$); (ii) prices are flexible and only wages are sticky ($\phi_W = 110$ and $\phi_P = 0$); (iii) $\alpha_\pi = 3$ in the model where both wages and prices are sticky; and (iv) the weight on the output gap in the

[15] The solutions to those problems are not necessarily REE because they are solved under the assumption that $\hat{C}_t$, $\hat{\pi}_t$, and $\hat{w}_t$ evolve as in equation 15, which is not consistent with an REE so long as $\mu_i \neq 0$ and $\varepsilon_{i,t} \neq 0$. However, conditional on those beliefs, we can use the same methods to compute the linear approximation to the solution to the model that we would use in the context when households had RE. When households and firms believe that μ_i is different from zero, that belief only affects the intercepts to the laws of motion in the solution to their problems. When households and firms believe that $\varepsilon_{i,t}$ can take values other than zero, those values appear linearly in the solution to the household an firm problems.

Taylor rule is zero ($\alpha_Y = 0$).

Table 1 reports that the value of b in the benchmark model is 0.62. This value implies that asymptotic convergence to the REE is very slow. According to Propositions 1 and 2, the asymptotic rate of convergence of $\|m_t\|$ is t^{b-1}. Using t^{b-1} as an approximation to $\|m_t\|$, the amount of time, $T_{2/3}$, that it takes to close two-thirds of a gap, $\|\hat{m}_0\|$ to the REE, is $T_{2/3} = (1/3)^{\frac{1}{b-1}}$ which corresponds to 20 quarters, i.e. or *5 years*. Table 1 also reports the amount of time, $\hat{T}_{2/3}$, that it takes to close two-thirds of the initial gap in beliefs in a learning equilibrium with initial beliefs as given the table. The values $\hat{T}_{2/3}$ show that asymptotic value, $T_{2/3}$, is a good approximation to the amount the small sample value, $\hat{T}_{2/3}$.

Table 1 also reports the eigenvectors associated with b, normalized to be of unit length and such that μ_π is non-negative. As explained in EH, to a first-order approximation, equation (1) can be written as

$$\theta_t \approx \theta_{t-1} + \gamma_t \left[D_1 M\left(0,0\right) - I \right] \theta_{t-1}. \tag{18}$$

Asymptotically, when θ_t is close to zero, the first-order approximation is particularly good.

Let V be the real matrix with columns composed of the real eigen vectors of $D_1 M\left(0,0\right)$ and the real and imaginary parts of the complex eigen vectors. The initial beliefs ($\hat{m}_0$) can be expressed so that

$$\hat{m}_0 = V \xi_0.$$

Here ξ_0 is a real vector of loadings on each eigenvalue. So long as the loading on the columns of V associated with the eigen value that has real part equal to b is not zero, beliefs will asymptotically converge at the rate t^{b-1}. To help provide intuition for why convergence in the simple NK is very slow, table 1 reports initial beliefs that are proportional to the column of V constructed from the eigenvectors associated with b. We use these initial beliefs because they focus attention on the implications of the eigen vectors associated with b. In the dynamic system given by equation 18, those eigenvectors converge to zero at the slowest rate. Focusing on other beliefs may give rise faster learning in simulations over short horizons, which could cloud intuition about why learning is asymptotically slow.

When b comes from a unique, non-complex eigenvalue of $D_1 M\left(0,0\right)$, the beliefs do not put a loading on any other column of V. While a loading on another column of V in the initial beliefs would have not any implication for the asymptotic rate of convergence, it could have meaningful implications for the equilibrium behavior of the model over short horizons which would cloud the intuition for slow asymptotic convergence. When b comes from a unique, complex-conjugate pair of eigenvalues of $D_1 M\left(0,0\right)$, then we focus on the real part of the eigenvectors associated with b.

When only wages or prices are sticky, people only have to form beliefs about aggregate

Table 1: b in the simple NK model

	b	$T_{2/3}$	$\hat{T}_{2/3}$	Initial beliefs: $[m_{0,C}, m_{0,\pi}, (m_{0,w})]$
Sticky prices and wages	0.63	20	25	$[-0.90, 0.39, -0.19]$
Only sticky prices	0.25	5	3	$[-0.93, 0.36]$
Only sticky wages	0.58	14	17	$[-0.92, 0.38]$
Higher α_π ($\alpha_\pi = 3$)	0.18	4	2	$[-0.98, 0.12, -0.15]$
Lower α_π ($\alpha_\pi = 1.01$)	0.98	very large	very large	$[-0.01, 1.00, -0.01]$
Lower α_Y ($\alpha_Y = 0$)	-0.30	3	1	$[1.00, 0.01, -0.03]$

Source: Authors' calculations.

consumption and inflation, so that the relevant vector of initial beliefs is two dimensional. When both wages and prices are sticky, people also have to form beliefs about the aggregate real wage and the vector of initial beliefs is three dimensional.

Below, we provide intuition for why beliefs converge to an REE and the nature of the forces governing the speed of convergence. In the discussion, unless stated otherwise, terms like "higher" and "lower" refer to the value of a variable relative to its magnitude in an REE. "More self-reinforcing beliefs" means that the current period ($t = 0$) value of a variable is closer to the current period expectations of the next period value of that variable. As a result, people will revise their expectations about the future value of the variable by a relatively small amount.

Consider the case when only prices are sticky. Table 1 reports that convergence occurs relatively quickly: $b = 0.25$ and $T_{2/3} = 5$ quarters. In this case, people initially expect the mean of consumption to be lower and the mean of inflation to be higher than in the REE. In addition, the deviation from the REE belief about consumption is larger than the deviation from the REE belief about inflation.

High expected inflation in subsequent periods leads firms to set higher prices in the initial period to avoid future adjustment costs. This behavior exerts *upwards* pressure on current period inflation. Because of adjustment costs, firms don't choose a current period firm-specific inflation rate as high as expected aggregate inflation.

Low expected aggregate consumption in the next period makes households feel poorer today because they will work less and receive less profits from monopolistic firms in the next period than they would in an REE. This negative wealth effect induces households to cutback on consumption and to supply more labor in the current period. This effect exerts *downward* pressure on inflation.

So, other things equal, current period consumption and the real wage rate are lower than their REE values. Because of competing forces, time zero inflation could be higher or lower than its REE value. Suppose, as is the case for our parameterized models, that the net effect of these forces is higher inflation. Then monetary policymakers would set nominal and real interest rates above their REE values. As long as monetary policy gives positive weight to the output gap, $\alpha_Y > 0$, there is a tradeoff between inflation and output objectives. Given the benchmark value of $\alpha_\pi(1.5)$ and $\alpha_Y(0.25)$, monetary policy reacts sufficiently aggressively to high inflation so that , so that the equilibrium value of inflation is substantially below expected inflation. Monetary policy also exerts downwards pressure on consumption. But initial beliefs about consumption were already very low. As it turns out time period zero aggregate consumption, is low but substantially above the expected value of time period 1.

Next period, people update their expectations based on current period data. The muted change in current period consumption and inflation, relative to what they expected inflation in period 1, induces people to *substantially* revise their expectations about future consumption and inflation. This process continues in future periods. The net result is a positive, relatively small value of b and quick convergence to the REE.

Consider next the case when only wages are sticky. Table 1 reports that convergence occurs more slowly than the sticky-price only case price: $b = 0.58$ and $T_{2/3}$ is larger than 3 years. The intuition is as follows. Initial beliefs are very close to their values in the sticky-price-only case. Also, as in the sticky-price-only case, low expected consumption makes households feel poorer so they reduce consumption. Absent price adjustment costs, firms set nominal price equal to nominal marginal cost which is equal to the nominal wage. In equilibrium, the real wage is equal to one, its REE value. So, in contrast to the sticky-price-only case, the equilibrium real wage rate is not lower than its REE value.

High expected inflation in subsequent periods leads households to set higher nominal wages in the initial period to avoid future adjustment costs. Because nominal marginal cost is also higher than in the sticky-price-only case, inflation is higher and more self-confirming. Other things equal, the higher inflation in the sticky wage only case results, through monetary policy, in a higher real interest rate. But that effect reinforces the household's desire to reduce consumption. On net, the equilibrium rate of inflation is higher and consumption falls by more than in the pure sticky price case. So the data is more self-confirming, b is higher and the rate of convergence will be slower.

Next consider the case in which both wages and prices are sticky. Here, the real wage need not be one. As in the previous cases, people initially expect the mean of consumption to be lower and the mean of inflation to be higher than in the REE. Because wages are sticky, the nominal wage cannot fall in the current period by as much as it did in the sticky-price-only case. Other things equal, this effect makes inflation higher than in the sticky-price-only case. Because prices are sticky, prices rise by more than in the sticky-wage-only case because adjustment costs induce firms to immediately raise prices in response to high expected inflation. Together, these effects cause inflation to be higher than in either the sticky-price-only or sticky-wage-only case (other things equal). Higher inflation induces monetary policy to raise the real interest rate by more than in those cases and there will be a larger fall in the equilibrium time period zero consumption relative to those cases. In equilibrium, inflation is higher than in either the sticky-price-only or sticky-wage only cases. Consequently, the equilibrium values of consumption and inflation are more self-fulfilling which leads to a higher value of b and a lower rate of convergence than in either of those two cases.

Next consider the case when α_π is higher than its benchmark value of 1.5. Relative to the benchmark case, the initial belief about consumption differs more from its REE value. The initial belief about inflation is closer to its REE value. A higher value of α_π means that monetary policy responds more aggressively to inflation. So, monetary policy offsets any rise in inflation induced by high inflation expectations by more so the real interest rate would be higher than if $\alpha_\pi = 1.5$. As a result, the equilibrium value of inflation is not very high relative to its REE and far from its expected time 1 value. Therefore, inflation beliefs are less self-reinforcing than in the benchmark case. Because inflation is low, the real interest rate and consumption don't change by much. So there is relatively large gap between equilibrium consumption and the initial expected value of time period one consumption. That, in turn, means that the initial belief about consumption, like inflation, is also less self-reinforcing. So b is smaller and the rate of convergence to the REE is faster.

Conversely, when α_π is lower than its benchmark value, learning is much slower. In this case, high expected inflation leads to high realized inflation and monetary policy does not raise the real interest rate by a meaningful amount. As a result, convergence of inflation and inflation beliefs back to their REE values is slow.

Finally, consider the case when $\alpha_Y = 0$. In this case, monetary policy only react to inflation. As Table 1 reports convergence is fast in this case. It is useful to consider what would happen if beliefs took the initial values we consider in our benchmark case. With those beliefs, monetary policy would react more aggressively to inflation. As a result, inflation and beliefs about inflation move to their REE values at a faster rate. With low expectations about aggregate output and the aggressive response to inflation, output is low, which causes inflation to overshoot its target value. As the monetary policy rule works to push inflation up, that also brings up output. Overall, the inflation and output gaps return to zero relatively quickly, which makes beliefs return to their REE values quickly.

5 Robustness

In this section we assess the robustness of our results to the following perturbations. First, we assess whether our basic results about the speed of convergence in the simple NK model hold in an empirically plausible, medium-size DSGE model. Second, to highlight the critical mechanisms that govern speed of convergence, we consider the speed of convergence in a different environment than the one considered in the previous section. The particular environment we focus on is one in which the ZLB is binding. Third, we consider the speed of convergence in a fully non-linear version of the simple NK model when the ZLB binds. Fourth, we consider the speed of convergence in the simple NK model when the ZLB binds when households and firms are internally rational in the sense of Adam and Marcet (2011).

5.1 An empirically plausible DSGE model

In this section we consider an empirically plausible DSGE model that is similar to the model in CEE. It shares the following features with CEE: sticky prices and wages, internal habit formation in consumption, endogenous rates of capital utilization, adjustment costs on investment, monetary policy is governed by a Taylor rule with interest rate smoothing.

The model differs in that we use Rotemberg-style adjustment costs rather than Calvo frictions to model nominal rigidities. Of course, in a linear REE context this difference is irrelevant. In appendix B, we provide a detailed specification of our variant of CEE. All parameters are the same as the point estimates reported by CEE. For convenience, these values are reported in appendix B. For the monetary policy rule, set the coefficients on the inflation rate and output gap to the values we used when analyzing the simple NK model. We set the interest rate smoothing parameter to 0.8, as in CEE.

As in the simple NK model, we assume that households and firms know how aggregate prices and quantities react to their lagged values and to structural shocks. Households and firms have to learn about the mean of those processes. As above, we calculate the speed of convergence of those beliefs. Table 2 reports the values of b and $T_{2/3}$ for seven versions of our model.

Table 2 reports that the speed of convergence to the REE in the benchmark model, labeled CEE, is very slow. It takes more than 60 years to close 2/3 of the gap between the initial beliefs and the REE beliefs. So moving from the simple NK model to an empirically plausible version of that model greatly exacerbates the time it takes converge to the REE. The intuition for the role of sticky prices and the response of monetary policy to the output gap (α_Y) is qualitatively similar to the intuition provided in the context of the simple NK model. As in the simple NK model, one key determinant of the speed of convergence to the REE occurs is sticky wages. Allowing wages to be flexible reduces $T_{2/3}$ from 243 to 27 (roughly 7 years).

Interest rate smoothing plays an important quantitative role in the speed of convergence. Without interest rate smoothing, the speed of convergence is actually somewhat faster than in the simple NK model. The intuition is straightforward. In the simple NK model (without interest rate smoothing), monetary policy moved aggressively to counter high inflation, albeit with some offset because of the output gap. With interest rate smoothing, monetary policy is no longer able to respond nearly as aggressively to counter high inflation. That, in turn, means that expectations about high inflation will be much more self-confirming.

Table 2 reports that habit formation plays very little role in determining the speed of convergence. The intuition is that even without habit formation concavity in the utility function is sufficient to prevent extremely large responses to initial beliefs. However, investment adjustment costs are very important. Absent those costs, it take thousands of years to close 2/3 of the gap to the REE. The intuition is straightforward. With no adjustment costs investment is very sensitive to initial beliefs, which makes movements in the output gap larger. Eliminating those gaps quickly would require tolerating higher inflation. When the Taylor rule puts weight on both objectives ($\alpha_Y > 0$), monetary policy will not induce large movements in either the inflation or the output gap that would lead people to significantly revise their initial beliefs.

As in the simple NK model, a key determinant of how quickly convergence to the REE occurs is sticky wages. When only prices are sticky, convergence is very fast. Moreover, as above, there is a positive interaction between sticky wages and prices in slowing down convergence to the REE. The intuition is exactly the same as in the simple NK model. Finally, we see that habit formation and adjustment costs in investment have a small effect on b and the speed of convergence.

5.2 Robustness of the mechanism to an alternative environment: the ZLB

In this subsection we highlight the critical mechanisms that govern speed of convergence by considering a different environment than the ones considered previously. The particular environment we focus on is one in which the ZLB is binding. To make the analysis as transparent as possible, we do that in the simple NK model. Surprisingly, we find that even that if we only allow for sticky prices, set $\alpha_Y = 0$, and have no interest rate smoothing convergence is very slow in that environment.

Up to the assumption of rational expectations, the environment is exactly the same as

Table 2: b in variants of the CEE (2005) model

	b	$T_{2/3}$	$\hat{T}_{2/3}$	Initial beliefs: $[m_{0,C}, m_{0,\pi}, m_{0,w}, m_{0,I}]$
CEE (2005)	0.80	243	345	$[-0.03, 0.99, 0.16, 0.04]$
No sticky wages	0.66	27	34	$[-0.09, 0.89, 0.44, 0.08]$
No sticky prices	0.77	125	173	$[-0.03, 0.97, 0, -0.23]$
No interest rate smoothing	0.60	16	19	[-0.37,0.74,-0.38,-0.41]
No habit formation	0.80	229	324	$[0.48, 0.86, 0.18, -0.03]$
No investment adjustment costs	0.86	2,173	3,160	$[-0.01, 0.05, 0.05, 1]$
No weight on output gap ($\alpha_Y = 0$)	0.02	4	2	$[-1, 0.08, -0.01, -0.03]$

Note: To implement no sticky wages, no sticky prices, and no adjustmnet costs on ivestment, we divide the benchmark values of the relevant adjustment costs parameters by 10,000. To implement no habit formation, we divide the benchmark value of γ by 10,000. To implement no interest rate smoothing, we divide the benchmark smoothing parameter by 10,000. To implement no weight on the output gap, we set $\alpha_Y = 0$.
Source: Authors' calculations.

the one considered in EW. To facilitate comparison with most of the related literature, we focus on the case of sticky prices.[16] We use the same notation is in the previous section. We begin by analyzing the properties of a linear approximation to the solution of the model.

In the current period, households discount next period's utility by $1/(1+r_t)$. In steady state, $r_t = r_{ss} > 0$. We assume that initially the economy is in the unique non-stochastic rational expectations steady state in which the nominal interest rate is positive. Then, unexpectedly, $r_t = r_\ell < r_{ss}$. People correctly understand that the next period's discount rate, r_{t+1}, is drawn from a two-state Markov chain, $r_{t+1} \in [r_\ell, r_{ss}]$, with an absorbing state:[17]

$$\Pr[r_{t+1} = r_\ell | r_t = r_\ell] = p, \;\; \Pr[r_{t+1} = r_{ss} | r_t = r_\ell] = 1 - p, \tag{19}$$
$$\Pr[r_{t+1} = r_\ell | r_t = r_{ss}] = 0.$$

Given our assumptions, under both RE and learning the only state variable is r_t and once $r_t = r_{ss}$ the economy returns to the initial rational expectations steady state.[18]

[16]An important exception are the papers by SGU that focus on downwardly rigid nominal wages.

[17]We have considered the scenario in which r_t is an ergodic process, i.e. when $r_t = r_{ss}$ there is a small, positive probability that discount rate shock and ZLB will occur again in the future. Our conclusions regarding the speed of learning are robust to this scenario.

[18]See Arifovic et al. (2018) for a discussion of stability for other learning models in which the deflationary steady state is learnable. There is another RE steady state in which there is deflation and the ZLB is binding (see Benhabib et al. (2001)). We abstract from that steady state equilibrium for now. Our conclusions are robust to the alternative assumption that households and firms know that they will return to the steady state emphasized by Benhabib et al. (2001).

5.2.1 The household and firm problems

Households' problem

The household value function when the ZLB is binding is given by:

$$V_h\left(b_{h,t-1},\Theta_t,x_t\right)=\max_{C_{h,t},b_{h,t}}\left\{\log\left(C_{h,t}\right)-\frac{\chi}{2}N_{h,t}^2\right. \tag{20}$$

$$\left.+\mathbb{E}_t\frac{p}{1+r_t}V_h\left(b_{h,t},\Theta_t,x_{t+1}\right)+\mathbb{E}_t\frac{1-p}{1+r_t}V_{h,ss}\left(b_{h,t}\right)\right\} \tag{21}$$

where

$$x_t=\begin{bmatrix}C_t\\ \pi_t\end{bmatrix}$$

subject to the budget constraint given in equation (12). Notice that the lagged real wage does not appear in x_t because wages are flexible. Also, $V_{h,ss}$ is the household value function when $r_t=r^{ss}$. The variable $b_{h,t}$ is the only argument in $V_{h,ss}$ because the household's value function is the same as under RE when $r_t=r^{ss}$.

We maintain the assumption that households solve their problem under anticipated utility. In addition, their information set, the way that they learn, and their priors about aggregate variables when $r_t=r^{\ell}$ are exactly the same as in the previous section.

Firms' problems

The technology for producing output and the market structure are exactly the same as in the previous section. The first order condition to the competitive firm's problem is still given by (3).

When $r_t=r^{\ell}$ the value function of the intermediate-good producing firm f is:

$$V_f\left(p_{f,t-1},\Theta_t,x_t\right)=\max_{p_{f,t}}\frac{1}{C_t}\left\{\left(p_{f,t}-\left(1-\tau_f\right)w_t\right)p_{f,t}^{-\varepsilon}Y_t-\frac{\phi}{2}\left(\frac{p_{f,t}}{p_{f,t-1}}\pi_t-1\right)^2C_t\right\}$$

$$+\mathbb{E}_t\frac{p}{1+r_t}V_f\left(p_{f,t},\Theta_t,x_{t+1}\right)+\mathbb{E}_t\frac{p}{1+r_t}V_{f,ss}\left(p_{f,t}\right). \tag{22}$$

where x_t is defined as in the household problem. $V_{f,ss}$ is the firm's value function when $r_t=r^{ss}$. The variable $p_{f,t}$ is the only argument in $V_{f,ss}$ because the firm's value function is the same as under RE when $r_t=r^{ss}$.

We maintain the assumption that firms solve their problem under anticipated utility. In addition, their information set, the way that they learn, and their priors about aggregate variables when $r_t=r^{\ell}$ are exactly the same as in the previous section.

5.2.2 Monetary and fiscal policy

Monetary policy sets the gross nominal interest rate, R_t, according to

$$R_t = \max\left\{1, 1 + r_{ss} + \alpha_\pi \left(\pi_t - 1\right)\right\}, \tag{23}$$

where $\alpha / \left(1 + r_{ss}\right) > 1$, and the max operator reflects the ZLB constraint.

For now, we assume that the level of , G_t, is equal to its non-stochastic steady-state value, $G_{ss} = 0.2 \times Y_{ss}$. As above, the government also pays a subsidy to intermediate goods firms, which it picks to offset steady state monopoly distortions. The government finances its expenditures with lump-sum taxes and balances its budget in each period.

The mapping from x_t to aggregate variables is the same as $F\left(x_t\right)$ in the previous section with the obvious modifications to account for government purchases.

We solve the model using a linear approximation around the REE (see the appendix for details). Finally the definition of a temporary equilibrium and a learning equilibrium is the same as in section 4.1.6.

5.2.3 The speed of convergence for the simple NK model in the ZLB

In what follows, we use the same parameter values as above (see section 4.3). In addition, we specify the two new parameters, p and r^ℓ, to be $p = 0.8$ and $r^\ell = -0.0015$.[19] In the $R > 1$ steady-state REE, $C_{ss} = 0.8$, $\pi_{ss} = 1$, $N_{ss} = 1$. In RE, the non-linear equilibrium is characterized by the number C^ℓ and π^ℓ. The model has multiple equilibria when $r_t = r^\ell$. Here, we focus on the unique equilibrium that is E-stable and linearize the model around that equilibrium.

Table 3: Eigenvalues of B

	b	Initial beliefs: $[m_{0,C}, m_{0,\pi}]$	$T_{2/3}$	$\widehat{T}_{2/3}$
ZLB	0.92	$[-0.98, -0.21]$	1.7 million	2.6 million

Table 3 also reports the value of b, the associated vector of initial beliefs, and the values of $T_{2/3}$ and $\hat{T}_{2/3}$. The initial beliefs correspond to the eigenvector associated with the eigenvalue b. It turns out that this vector is almost exactly proportional to the vector $\left[\log\left(C_\ell / C_{ss}\right), \log\left(\pi_\ell / \pi_{ss}\right)\right]$. A natural starting point for beliefs is that the belief that con-

[19] It is well known that the REE is sensitive to different values of p. We find that b is an increasing function of p. The value $p = 0.76$ is approximately the smallest value of p for which the ZLB binds in our model. When $p = 0.76$, $b = 0.81$. Larger values of p result in even slower convergence.

sumption and inflation are going to be equal to their REE steady state values. The set of initial beliefs reported in Table 3 is essentially the same as that natural starting point.

Our key result here is that the learning economy converges very slowly when the ZLB is binding: $T_{2/3} \approx$ hundreds of thousands of years. The reason is that absent a policy response to inflation in the ZLB beliefs about expected inflation are very self-confirming. To understand why, suppose people's prior is that inflation will be high in the next period, causing firms to want to raise prices in the current period. When the ZLB, the central bank cannot take action in the current period to lower actual inflation. Consequently, expectations are more self-fulfilling and people only adjust their beliefs very slowly.

In the benchmark small NK model, our benchmark variant of the CEE model, and in the model with a binding a binding ZLB considered here, learning is very slow. Moreover, the basic intuition about almost self-confirming beliefs applies to all of those models. In this sense, our results are robust to considering learning in alternative environments.

5.3 Speed of convergence in a fully non-linear model

In this subsection we consider the robustness of our learning results to working with a fully non-linear solution of the simple NK model. Our most dramatic results about the speed of convergence arise when the ZLB is binding. So, we work with the simple NK model when the ZLB is binding. Appendix provides a detailed description of how we solved the model. For now, we solve the model under anticipated utility.

Similar to above, we assume that households and firms form beliefs about $x_t = [C_t, \pi_t]$. People believe

$$\widehat{x}_t = \begin{bmatrix} \widehat{C}_t \\ \widehat{\pi}_t \end{bmatrix} = \begin{bmatrix} \mu_C \\ \mu_\pi \end{bmatrix} + \begin{bmatrix} \varepsilon_{C,t} \\ \varepsilon_{\pi,t} \end{bmatrix}$$

where $\varepsilon_{C,t} \sim N(0, \sigma_C^2)$ and $\varepsilon_{\pi,t} \sim N(0, \sigma_\pi^2)$. These distributions are independent across time. People are uncertain about the values of μ_i, σ_i^2 for $i \in \{C, \pi\}$. In contrast to the linear approximation to the non-linear model, people's beliefs about the variance of shocks matters.

Section 4.2 contains a complete description of people's priors about about μ_i and σ_i^2. That section also contains a description of the associated posterior distributions. Because we are working with the non-linear solution to the model, we need to specify people's initial beliefs for both μ_i and σ_i^2. We set these initial values so that

$$\begin{pmatrix} m_C, & m_\pi, & 1/\lambda_C, & 1/\lambda_\pi, & \psi_C, & \psi_\pi, & \alpha_C, & \alpha_\pi \end{pmatrix}' =$$
$$\begin{pmatrix} \log(C_{ss}), & \log(\pi_{ss}), & 1, & 1, & 0.02, & 0.02, & 1/2, & 1/2 \end{pmatrix}'$$

We set $m_C = \log(C_{ss})$ and $m_\pi = \log(\pi_{ss})$ because, as discussed above, this set of beliefs is a natural starting point and is essentially the same as beliefs based on the eigenvector associated with the maximal eigenvalue of the model, b, reported in Table 3.

The solid and dashed lines in figure 1a display the evolution of inflation, consumption, and the real interest rate after the drop in r under REE and learning, respectively. In this figure, we assume that $r_t = r_\ell$ throughout the period shown. Of course, after a few quarters, this is a very low probability event. Nevertheless, it is useful to consider how the economy would evolve in this event to analyze the speed of convergence. Two key features are worth noting. First, the outcomes for consumption and inflation are very different in the REE than in the learning equilibrium. In the REE, there is a very large drop in inflation and consumption, and the real interest rate rises sharply. The fall in inflation and consumption and the rise in the real rate are much smaller under learning. Second, the learning economy converges very slowly to the REE. After people initially change their views somewhat quickly, the rate at which they change their views slows dramatically (see figure 1b).

Figure 1: Learning equilibrium in non-linear model

(a) Outcomes

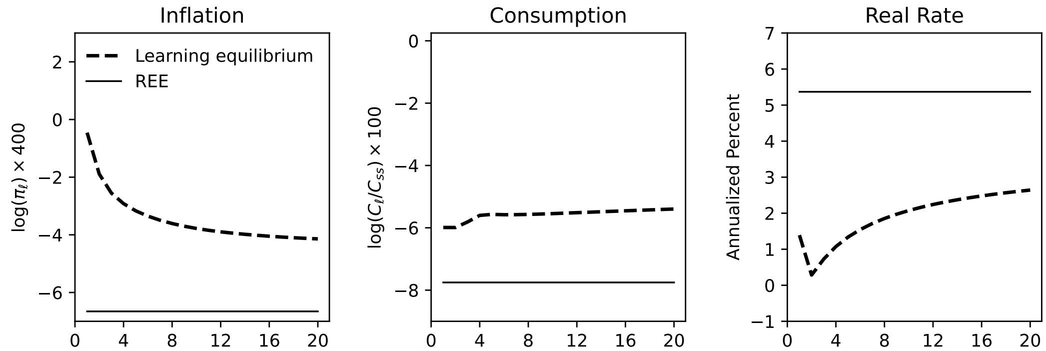

(b) Beliefs

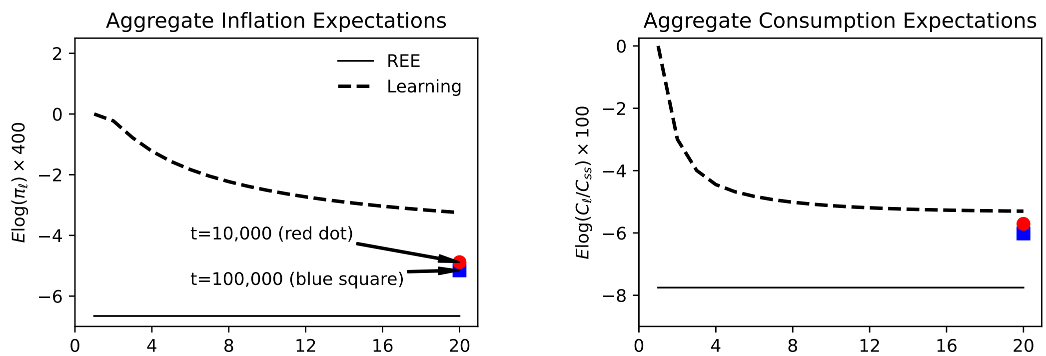

The dot labeled $t = 10,000$ in figure 1b displays people's views about the variables after 10,000 quarters. Given our value, $p = 0.8$, r is only expected to be low for about five quarters. The probability of staying at the ZLB for $10,000$ quarters is $p^{10,000} \approx 0$. So, the fact that the system converges to the REE is essentially irrelevant because the ZLB is likely to be over well before that time. The crucial point is that in a typical ZLB episode, the associated REE is very different from the learning equilibrium for an enormous amount of time.

Finally, our results about slow convergence when the ZLB is binding are consistent with those that we obtain with the linear approximation to the model's solution.

5.4 Speed of convergence with internalized learning in a fully non-linear model

In this subsection, we consider the robustness of our result to the assumption that people make decisions under anticipated utility. For comparability with the results in the previous subsection, we again work with the simple NK model when the ZLB is binding.

Recall that anticipated utility differs from internalized learning in three ways. First, under anticipated utility people do not update their beliefs based on current-period data until after they have made their decisions in the current period. Second, under anticipated utility people ignore that in the future when they see more data they will update their views about unknown parameters. Third, under anticipated utility people ignore their uncertainty about the mean and variance of unknown parameters. Under internalized learning people update their beliefs in response to the information they see in the current period and fully integrate the fact that they are learning when they solve their problems. See appendix C for a description of our computational approach to solving the model under internalized learning.

Figure 2 compares the evolution of the learning equilibrium under anticipated utility (dashed black line) and internalized learning (dash-dot blue line). The figure also includes the REE outcomes. The key takeaway is that we obtain the same slow-learning result regardless of which approach we take to decision making. However, consumption and inflation fall somewhat more under internalized learning. The reason is that under anticipated utility people do not update their beliefs in the initial period in response to low consumption and inflation until after they have made their decisions. Under internalized learning, people update their beliefs before they make their decisions, which causes their beliefs about future inflation and consumption to be lower. In any event, consumption, inflation, and the real interest rate are very different than the REE outcomes.

Figure 2: Anticipated utility versus internalized learning

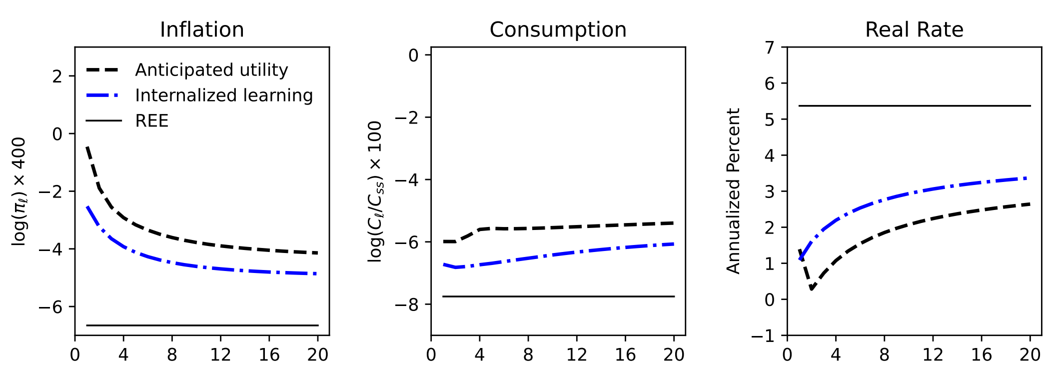

6 Does learning matter for policy?

In this section, we address the question of how learning affects the implications of monetary and fiscal policy. It is well known that policy implications of the ZLB are very dramatic.

For example, Eggertsson (2010) and Christiano et al. (2011) argued that the size of the government purchases multiplier is much larger when the ZLB is binding. Additionally, Del Negro et al. (2023) and other have argued that forward guidance about the nominal interest rate is particularly effective when the ZLB is binding. They also emphasize the forward guidance puzzle, according to which a promise to change the nominal interest rate far in the future can have nearly the same effect as changing the nominal interest rate in the current period. Because of these dramatic implications for policy, we work with the non-linear solution to the simple NK model when the ZLB is binding and people internalize learning in their decision making. As it turn out, similar conclusions obtain if we work with anticipated utility or the linear approximation to the solution of the model.

6.1 The government purchases multiplier

In this subsection we compare the efficacy of fiscal policy under rational expectations and internalized learning. We compute the government purchases multiplier by considering the effect on GDP, $C_t + G_t$, of a 5 percent rise in government purchases relative to its steady-state level, $G_t(r_\ell) = 1.05 \times G_{ss}$. We define the multiplier in the ZLB at time $t+j$ as

$$\frac{\Delta C_{t+j} + \Delta G_{t+j}}{\Delta G_{t+j}}. \tag{24}$$

Here, $\Delta G_{t+j} = 0.05 \times G(r_{ss})$ and ΔC_{t+j} is the difference between consumption in the equilibrium when G_{t+j} is high relative to its value when G_{t+j} is equal to its steady state value.

In the REE, the multiplier in the ZLB for all $j \geq 0$ is equal to 3.95 (see figure 3). The multiplier is large when the ZLB is binding because the rise in G generates an increase in inflation and then expected inflation. Because R_t is fixed, this rise generates a fall in the real interest rate and a rise in C_t. So, in this case, the multiplier is bigger than one.

Figure 3: Equilibria with and without an increase in G_t

(a) Outcomes during ZLB

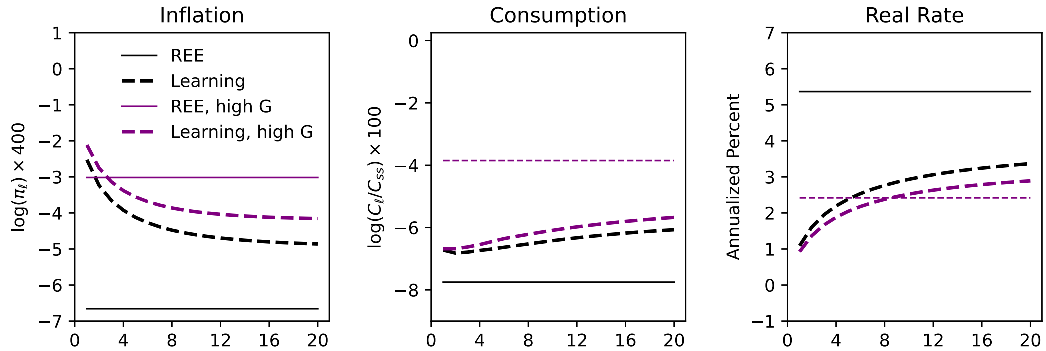

(b) Government purchases multiplier during ZLB

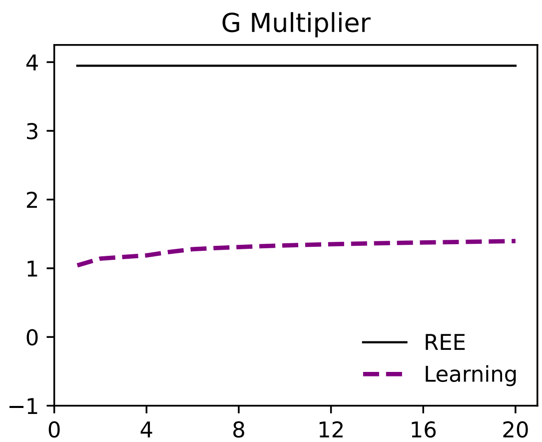

Under learning, expected inflation is partially backward-looking and does not move much with a rise in G_t. As a result, the real interest rate does not fall very much, and the response in consumption is small. Figure 3b displays the value of the multiplier over time in the REE and under learning in the ZLB. We conclude that under learning the effect of an increase in government purchases is small compared to its effect in the REE over the 20 quarters displayed.

6.2 Forward guidance about R_t

Under forward guidance, the monetary authority commits to keeping the nominal interest rate at the ZLB for J periods after r_t has returned to its steady-state level. To make our point as simply as possible, we consider the case $J = 1$. In the appendix we show that the number of REE proliferates under forward guidance. However, only one of those equilibria is stable under learning. We focus on this equilibrium and analyze the effect of forward guidance. The appendix also contains all computational details for the learning equilibria with forward guidance.

In the period of forward guidance, $r_t = r_{ss}$, $R_t = 1$. In all periods when $r_t = r_\ell$, people understand that the economy reverts to an REE when $r_t = r_{ss}$. We consider the learning equilibrium using the same initial values for Θ_t as in section 5.3. The learning equilibrium under forward guidance is indistinguishable from the learning equilibrium without forward guidance, as shown in Figure 4. The reason is that forward guidance has small effects on inflation and consumption in period the period when r_t reverts to r_{ss}. In turn, these small effects have almost no effect on the learning equilibrium where beliefs are partially backward looking. Clearly, there is no forward guidance puzzle under learning.[20] That puzzle emerges under RE because of the strong effect of forward guidance on expected inflation. Under learning, expectations are backward-looking, and forward guidance has little influence on expected inflation. Note that the learning equilibria with and without forward guidance converge slowly.

Figure 4: Forward guidance under learning and in the REE

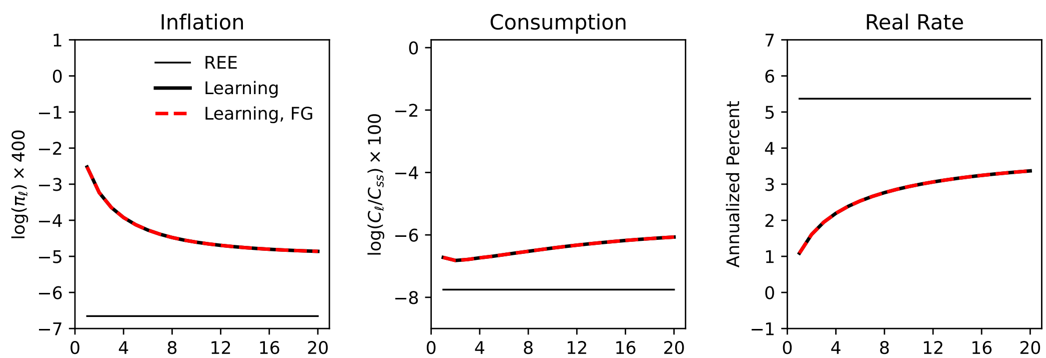

7 Conclusion

In this paper, we consider the speed with which people learn about their environment. To characterize the speed of convergence of people's beliefs, we analytically extend results in the literature to characterize the asymptotic convergence rate of multivariate systems. We argue that the slow convergence result emerges naturally in empirically plausible NK models, including cannonical DSGE models that are similar to CEE (2005). The basic reason is that expectations are partially self-fulfilling in these models. When learning is slow, analyses of fiscal and monetary policies under rational expectations can be very misleading.

[20] For a discussion of the forward guidance puzzle, see Del Negro et al. (2023). Under anticipated utility (not displayed) forward guidance has a slight effect, but not large enough to be economically meaningful.

Finally, we note that there are other circumstances in which slow learning could arise. For example, plausible parameterizations of Cagan (1956)'s model of money demand under hyperinflation map into high b economies. Results in Marcet and Sargent (1995, Table 6.3) imply that estimates of the elasticity of money demand in hyperinflations (see, for example, Christiano (1987) and Taylor (1991)) map into high values of b.

References

Adam, Klaus, Albert Marcet, and Johannes Beutel, "Stock price booms and expected capital gains," *American Economic Review*, 2017, *107* (8), 2352–2408.

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

_ **and Sebastian Merkel**, "Stock price cycles and business cycles," *European Central Bank Working Paper no. 2316*, September 2019.

Angeletos, George-Marios and Chen Lian, "Dampening general equilibrium: From micro to macro," Technical Report, National Bureau of Economic Research 2017.

_ **and** _ , "Forward Guidance without Common Knowledge," *American Economic Review*, September 2018, *108* (9), 2477–2512.

Arifovic, Jasmina, Stephanie Schmitt-Grohé, and Martin Uribe, "Learning to live in a liquidity trap," *Journal of Economic Dynamics and Control*, 2018, *89*, 120–136.

Benhabib, Jess, Stephanie Schmitt-Grohé, and Martin Uribe, "Monetary policy and multiple equilibria," *American Economic Review*, 2001, *74* (2), 165–170.

Cagan, Phillip, "The monetary dynamics of hyperinflation," *Studies in the Quantity Theory if Money*, 1956.

Christiano, Lawrence J, "Cagan's model of hyperinflation under rational expectations," *International Economic Review*, 1987, pp. 33–49.

Christiano, Lawrence J., Martin S. Eichenbaum, and Sergio Rebelo, "When is the Government Spending Multiplier Large?," *Journal of Political Economy*, February 2011.

Christopeit, Norbert and Michael Massmann, "Estimating structural parameters in regression models with adaptive learning," *Econometric Theory*, 2018, *34* (1), 68–111.

Cogley, Timothy and Thomas J Sargent, "Anticipated utility and rational expectations as approximations of Bayesian decision making," *International Economic Review*, 2008, *49* (1), 185–221.

Eggertsson, Gauti B., "What Fiscal Policy Is Effective at Zero Interest Rates?," in "NBER Macroeconomic Annual," Vol. 25, University of Chicago Press, 2010, pp. 59–112.

Eggertsson, Gauti B and Michael Woodford, "Policy options in a liquidity trap," *American Economic Review*, 2004, *94* (2), 76–79.

Erceg, Christopher J. and Andrew T. Levin, "Imperfect credibility and inflation persistence," *Journal of Monetary Economics*, 2003, *50* (4), 915–944.

_ **, Dale W. Henderson, and Andrew T. Levin**, "Optimal monetary policy with staggered wage and price contracts," *Journal of Monetary Economics*, 2000, *46* (2), 281–313.

Eusepi, Stefano, Christopher Gibbs, and Bruce Preston, "Monetary policy trade-offs at the zero lower bound," *Available at SSRN 4069078*, 2022.

Evans, George W, "Theories of learning and economic policy," *Revue d'économie politique*, 2021, *132* (3-4), 583–608.

_ and Seppo Honkapohja, "Local convergence of recursive learning to steady states and cycles in stochastic nonlinear models," *Econometrica*, 1995, pp. 195–206.

_ and _ , "Convergence for difference equations with vanishing time-dependence, with applications to adaptive learning," *Economic Theory*, 2000, *15*, 717–725.

Evans, George W. and Seppo Honkapohja, *Learning and Expectations in Macroeconomics*, Princeton University Press, 2001.

Evans, George W and Seppo Honkapohja, "Monetary policy, expectations and commitment," *Scandinavian Journal of Economics*, 2006, *108* (1), 15–38.

Farhi, E, M Petri, and I Werning, "The fiscal multiplier puzzle: Liquidity traps, bounded rationality, and incomplete markets," Technical Report, mimeo 2020.

Farhi, Emmanuel and Iván Werning, "Monetary policy, bounded rationality, and incomplete markets," *American Economic Review*, 2019, *109* (11), 3887–3928.

Farmer, Leland, Emi Nakamura, and Jón Steinsson, "Learning about the long run," Technical Report, National Bureau of Economic Research 2021.

Ferrero, Giuseppe, "Monetary policy, learning and the speed of convergence," *Journal of Economic Dynamics and Control*, 2007, *31* (9), 3006–3041.

García-Schmidt, Mariana and Michael Woodford, "Are low interest rates deflationary? A paradox of perfect-foresight analysis," *American Economic Review*, 2019, *109* (1), 86–120.

Gust, Christopher, Edward Herbst, and David López-Salido, "Forward guidance with bayesian learning and estimation," *FEDS Working Paper*, 2018.

Heemeijer, Peter, Cars Hommes, Joep Sonnemans, and Jan Tuinstra, "Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation," *Journal of Economic dynamics and control*, 2009, *33* (5), 1052–1072.

Hirsch, Morris W. and Stephen Smale, *Differential equations, dynamical systems, and linear algebra*, Academic Press, 1974.

Hommes, Cars, "The heterogeneous expectations hypothesis: Some evidence from the lab," *Journal of Economic dynamics and control*, 2011, *35* (1), 1–24.

Horn, R. A. and C. A. Johnson, *Matrix Analysis*, Cambridge University Press, 1985.

Iovino, Luigi and Dmitriy Sergeyev, "Central bank balance sheet policies without rational expectations," *Review of Economic Studies*, 2023, *90* (6), 3119–3152.

Kaplan, Greg and Giovanni L Violante, "Microeconomic heterogeneity and macroeconomic shocks," *Journal of Economic Perspectives*, 2018, *32* (3), 167–94.

Kreps, David M, "Anticipated utility and dynamic choice," *Econometric Society Monographs*, 1998, *29*, 242–274.

Ljung, Lennart, "Analysis of recursive stochastic algorithms," *IEEE transactions on automatic control*, 1977, *22* (4), 551–575.

Marcet, Albert and Thomas J Sargent, "Convergence of least-squares learning in environments with hidden state variables and private information," *Journal of Political Economy*, 1989, *97* (6), 1306–1322.

_ **and** _ , "Convergence of least squares learning mechanisms in self-referential linear stochastic models," *Journal of Economic theory*, 1989, *48* (2), 337–368.

_ **and Thomas Sargent**, "Speed of Convergence of Recursive Least Squares: Learning with Autoregressive Moving-Average Perceptions," in Alan Kirman and Mark Salmon, eds., *Learning and rationality in economics*, Oxford University Press, UK, 1995, chapter 6, pp. 179–215.

Negro, Marco Del, Marc P Giannoni, and Christina Patterson, "The forward guidance puzzle," *Journal of Political Economy Macroeconomics*, 2023, *1* (1), 43–79.

Preston, Bruce, "Adaptive learning in infinite horizon decision problems," *Unpublished, Columbia University*, 2005.

Rotemberg, Julio J, "Sticky prices in the United States," *Journal of political economy*, 1982, *90* (6), 1187–1211.

Schmitt-Grohé, Stephanie and Martın Uribe, "Closing small open economy models," *Journal of international Economics*, 2003, *61* (1), 163–185.

Sims, Christopher A., "Solving Linear Rational Expectations Models," *Computational Economics*, 2001, *20*, 1–20.

Taylor, John B., "A Historical Analysis of Monetary Policy Rules," in "Monetary Policy Rules," University of Chicago Press, 1999, pp. 319–341.

Taylor, Mark P, "The hyperinflation model of money demand revisited," *Journal of money, Credit and Banking*, 1991, *23* (3), 327–351.

Vimercati, Riccardo Bianchi, Martin S Eichenbaum, and Joao Guerreiro, "Fiscal policy at the zero lower bound without rational expectations," Technical Report, National Bureau of Economic Research 2021.

Vives, Xavier, "How fast do rational agents learn?," *The Review of Economic Studies*, 1993, *60* (2), 329–347.

Woodford, Michael, "Self-fulfilling expectations and fluctuations in aggregate demand," 1990.

_ , "Methods of Policy Accommodation at the Interest-rate Lower Bound," *Proceedings-Economic Policy Symposium-Jackson Hole, Federal Reserve Bank of Kansas City*, 2012, pp. 185–288.

_ **and Yinxi Xie**, "Policy options at the zero lower bound when foresight is limited," in "AEA Papers and Proceedings," Vol. 109 2019, pp. 433–437.

_ **and** _ , "Fiscal and monetary stabilization policy at the zero lower bound: Consequences of limited foresight," *Journal of Monetary Economics*, 2022, *125*, 18–35.

A Proofs of Propositions 1 and 2

The basic strategy of the proofs of Propositions 1 and 2 is similar to the strategy of the proofs of Propositions 1 and 2 in Evans and Honkapohja (2000).

A.1 Proposition 1

Proof. As in Evans and Honkapohja (2000), write equation 1 as

$$\theta_t = (I + \gamma_t J)\,\theta_{t-1} + \gamma_t r_t$$

where $J = D_1 M(0,0) - I$ and the equation defines r_t. From Proposition 1 in Evans and Honkapohja (2000), there exists neighborhood U_1 such that for every $\theta_0 \in U_1$, θ_t converges to 0.

Let $1 + \gamma_t \lambda$ be an eigenvalue of $I + \gamma_t J$. Here, $\lambda = \lambda_r + i\lambda_c$ is an eigenvalue of J . Fix a value of $\delta > 0$. We now show that for t large enough, all of the eigenvalues of $I + \gamma_t J$ are within a circle of radius $|1 + \gamma_t \lambda| < 1 + \gamma_t(b - 1 + \delta/4)$. Note that

$$\begin{aligned}|1 + \gamma_t \lambda| &= \sqrt{[1 + \gamma_t \lambda_r]^2 + (\gamma_t \lambda_c)^2} \\ &= \sqrt{1 + 2\gamma_t(\lambda_r + \delta/4) - 2\gamma_t \delta/4 + \gamma_t^2 \lambda_r^2 + \gamma_t^2 \lambda_c^2}.\end{aligned}$$

For large enough t,

$$0 < 1 + 2\gamma_t(\lambda_r + \delta/4) \leq 1 + 2\gamma_t(b - 1 + \delta/4),$$

where the first inequality requires t sufficiently large and the second inequality is true by definition of b. Additionally, for large enough t,

$$-2\gamma_t\delta/4 + \gamma_t^2\lambda_r^2 + \gamma_t^2\lambda_c^2 < 0 < \gamma_t^2\left(b - 1 + \delta/4\right)^2.$$

So, for large enough t,

$$|1 + \gamma_t\lambda| < \sqrt{1 + 2\gamma_t\left(b - 1 + \delta/4\right) + \gamma_t^2\left(b - 1 + \delta/4\right)^2},$$

which immediately implies $|1 + \gamma_t\lambda| < 1 + \gamma_t(b - 1 + \delta/4)$.

As in step II of the proof of Proposition 5.2 in Evans and Honkapohja (1995), who follow Horn and Johnson (1985), there exists a norm so that for large enough t,

$$\|I + \gamma_t J\| \leq 1 + \gamma_t\left(b - 1 + \frac{1}{2}\delta\right).$$

Now consider r_t. As in Evans and Honkapohja (1995), the maintained assumptions imply local Lipschitz continuity and that for large enough t (small enough γ_t) there exists a neighborhood of the origin U_2, where $U_2 \subset U_1$,

$$\|r_t\| \leq \frac{\delta}{4}\|\theta_{t-1}\|.$$

It follows from the sub-additivity and sub-multiplicativity of the matrix and vector norms that for $\theta_{t-1} \in U_2$ and for large enough t

$$\begin{aligned}\|\theta_t\| &\leq \|I + \gamma_t J\|\,\|\theta_{t-1}\| + \frac{\delta}{4}\|\theta_{t-1}\|\gamma_t \\ &\leq \left(1 + \gamma_t\left(b - 1 + \frac{3}{4}\delta\right)\right)\|\theta_{t-1}\|.\end{aligned}$$

Consider $\theta_0 \in U_2$. For $t - 1 \geq 1$, define $\tilde{\theta}_{t-1} = \frac{\theta_{t-1}}{(t-1)^{b-1+\delta}}$. Then for large enough t there exists a finite, positive constant g so that

$$\begin{aligned}\left\|\tilde{\theta}_t\right\| &\leq \left(1 + \frac{1}{t-1}\right)^{1-b-\delta}\left(1 + \gamma_t\left(b - 1 + \frac{3}{4}\delta\right)\right)\left\|\tilde{\theta}_{t-1}\right\| \\ &\leq \left(1 + \frac{1-b-\delta}{t-1} + g\left(\frac{1}{t-1}\right)^2\right)\left(1 + \gamma_t\left(b - 1 + \frac{3}{4}\delta\right)\right)\left\|\tilde{\theta}_{t-1}\right\|\end{aligned}$$

The first line follows from absolute homogeneity of a matrix norm. The second line follows by from Taylor's theorem.

The strategy in what follows is similar to the one used in Evans and Honkapohja (2000), except that we focus on $\tilde{\theta}_t$ rather than θ_t. The inequality above can be written as

$$\left\|\tilde{\theta}_t\right\| \leq \tilde{k}_t \left\|\tilde{\theta}_{t-1}\right\|$$

where

$$\begin{aligned}\tilde{k}_t =& 1 + \frac{1-b-\delta}{t-1} + \gamma_t \left(b - 1 + \frac{3}{4}\delta\right) \\ &+ \frac{1-b-\delta}{t-1}\gamma_t \left(b-1+\frac{3}{4}\delta\right) + g\left(\frac{1}{t-1}\right)^2 \left(1 + \gamma_t \left(b - 1 + \frac{3}{4}\delta\right)\right).\end{aligned}$$

For t large enough $\tilde{k}_t < 1$ and for all $\left\|\tilde{\theta}_{t-1}\right\| \neq 0$

$$\left\|\tilde{\theta}_t\right\| - \left\|\tilde{\theta}_{t-1}\right\| \leq \left(\tilde{k}_t - 1\right) \left\|\tilde{\theta}_{t-1}\right\| < 0. \tag{25}$$

Thus, there exists a τ so that if $t > \tau$ then $\left\|\tilde{\theta}_t\right\|$ is a decreasing. Let $t - 1 > \tau$. Following the strategy in Evans and Honkapohja (2000), we iterate equation 25 and obtain

$$1 - \frac{\left\|\tilde{\theta}_{t-1}\right\|}{\left\|\tilde{\theta}_{t+s}\right\|} \leq \sum_{i=t}^{t+s} \left(\tilde{k}_i - 1\right).$$

Note that

$$\lim_{s\to\infty} \sum_{i=t}^{t+s} \frac{1-b-\delta}{i-1}\gamma_i \left(b - 1 + \frac{3}{4}\delta\right) < \infty,$$

$$\lim_{s\to\infty} \sum_{i=t}^{t+s} g\left(\frac{1}{i-1}\right)^2 \left(1 + \gamma_i \left(b - 1 + \frac{3}{4}\delta\right)\right) < \infty.$$

Additionally, for large enough i there exists a $\tilde{\omega} > 0$ so that

$$\frac{1-b-\delta}{i-1} + \gamma_i \left(b - 1 + \frac{3}{4}\delta\right) < -\frac{1}{i}\tilde{\omega}.$$

So,

$$\lim_{s\to\infty} \sum_{i=t}^{t+s} \left(\tilde{k}_i - 1\right) = -\infty,$$

meaning $\lim_{t\to\infty} \left\|\tilde{\theta}_t\right\| = 0.$ □

A.2 Proposition 2

Proof. From Proposition 1 in Evans and Honkapohja (2000), there exists $\kappa_1 > 0$ so that if $\theta_0 \subset B_{\kappa_1}(0)$,where $B_{\kappa_1}(0)$ is an open ball around the origin, then θ_t converges. We will only consider values of $\theta_0 \subset B_{\kappa_1}(0)$.

As in Evans and Honkapohja (2000), define $J \equiv D_1 M(0,0) - I$ and write equation 1 as

$$\theta_t = \theta_{t-1} + \gamma_t \left[J\theta_{t-1} + r_t\right]$$

where r_t is the residual. Use the real and imaginary parts of the generalized eigenvectors of J to form a basis for $\mathbb{R}^k$ (those vectors are a basis that put J into real Jordan form). Define E_A to be the span of the real and imaginary parts of the generalized eigenvectors associated with eigenvalues with real parts less than $b-1$ and E_B to be the span of the real and imaginary parts of the generalized eigenvectors associated with eigenvalues with real parts equal to $b-1$.[21] Note that $\mathbb{R}^k = E_A \oplus E_B$. Define $x_t, r_{A,t} \in E_A$ and $y_t, r_{B,t} \in E_B$ to be the unique vectors so that

$$\begin{aligned} \theta_t &= x_t + y_t \\ r_t &= r_{A,t} + r_{B,t}. \end{aligned}$$

It is useful to note that there are matrices $A = J|_{E_A}$ and $B = J|_{E_B}$ so that

$$M(\theta_{t-1}, \gamma_t) - \theta_{t-1} = Ax_{t-1} + r_{A,t} + By_{t-1} + r_{B,t},$$

and that the eigenvalues of A are the eigenvalues of J with real parts less than $b-1$, and that the eigenvalues of B are the eigenvalues of J with real parts equal to $b-1$. Let $s-1$ be the smallest real part of eigenvalues of J and let $a-1$ (where $a-1 < b-1$) be the second-largest real part. Following Evans and Honkapohja (2000) by applying the lemma on page 145 of Hirsch and Smale (1974), for any $\delta_a, \delta_b > 0$, there exists Euclidean norms on E_A and E_B so that

$$\begin{aligned} &(s-1-\delta_a)\|x\|^2 \leq \langle \gamma_t Ax, x\rangle \leq (a-1+\delta_a)\|x\|^2 \ \forall x \in E_A. \\ &(b-1-\delta_b)\|y\|^2 \leq \langle \gamma_t By, y\rangle \ \forall y \in E_B. \end{aligned}$$

[21]To ease exposition, we assume that there are eignevalues with real part less than $b-1$. If not, the proof would proceed with straightforward modification. Note that proposition 2 in Evans and Honkapohja (2000) considers the case when $b > 1$, while we consider the case when $b < 1$. Evans and Honkapohja (2000) partition the space by using eigenvectors associated with positive or negative real parts. We partition the space by using the eigenvectors with real parts equal to $b-1$ and those with real parts less than $b-1$.

Choose δ_a and δ_b so that $a + \delta_a < b - \delta_b$ and $2\delta_b < \delta$, where δ is given in the statement of the proposition. As in Evans and Honkapohja (2000), define $\|\theta\| = \sqrt{\|x\|^2 + \|y\|^2}$.

As in Evans and Honkapohja (2000) for any $\epsilon > 0$, there exists a δ_r so that if $\theta_{t-1} \in B_{\delta_r}(0)$

$$\|r_t\| \leq \epsilon \|\theta_{t-1}\|.$$

Note that if $\theta_{t-1} \in B_{\delta_r}(0)$

$$\langle Ax_{t-1} + r_{A,t}, x_{t-1} \rangle \leq a - 1 + \delta_a + \epsilon\delta_r$$

and

$$\langle By_{t-1} + r_{B,t}, y_{t-1} \rangle \geq b - 1 - \delta_b - \epsilon\delta_r.$$

Similar to Evans and Honkapohja (2000), for $\beta > 0$ define

$$C_\beta = \{\theta : \theta = x + y, x \in E_A, y \in E_B | \|y\| \geq \beta \|x\|\}$$

and note that for $\theta_t \in C_\beta \cap B_{\delta_r}(0)$, the Cauchy-Schwartz inequality implies

$$\begin{aligned}
&\langle \gamma_{t+1} Ax_t + \gamma_{t+1} By_t + \gamma_{t+1} r_{t+1}, x_t + y_t \rangle \\
&= \gamma_{t+1} \langle Ax_t, x_t \rangle + \gamma_{t+1} \langle By_t, y_t \rangle + \gamma_{t+1} \langle r_{t+1}, \theta_t \rangle \\
&\geq \gamma_{t+1} (s - 1 - \delta_a) \|x_t\|^2 + \gamma_{t+1} (b - 1 - \delta_b) \|y_t\|^2 - \gamma_{t+1}\epsilon \|\theta_t\|^2 \\
&\geq \gamma_{t+1} \frac{(s - 1 - \delta_a)}{\beta} \|y_t\|^2 + \gamma_{t+1} (b - 1 - \delta_b) \|y_t\|^2 - \gamma_{t+1}\epsilon \|\theta_t\|^2 \\
&\geq \gamma_{t+1} \left[\frac{(s - 1 - \delta_a)}{\beta} + (b - 1 - \delta_b) - \epsilon \right] \|\theta_t\|^2.
\end{aligned}$$

Select β so that for small enough ϵ

$$\langle \gamma_{t+1} Ax_t + \gamma_{t+1} By_t + \gamma_{t+1} r_{t+1}, x_t + y_t \rangle \geq \gamma_{t+1} (b - 1 - 2\delta_b) \|\theta_t\|^2.$$

Following the strategy in Evans and Honkapohja (2000), we now show that for large enough t if $\theta_{t-1} \in C_\beta \cap B_{\delta_r}(0)$ then $\theta_t \in C_\beta \cap B_{\delta_r}(0)$. Because θ_t converges, we can select τ large enough so that for $t \geq \tau$, $\theta_t \in B_{\delta_r}(0)$. So, we only need to show that, for large enough t, if $\theta_{t-1} \in C_\beta$ and $\|y_{t-1}\| \neq 0$ then $\theta_t \in C_\beta$. The steps in the proof of Proposition 2 in Evans and Honkapohja (2000) related to g_β establish this result for large enough t and small enough ϵ.[22] Then, for τ large enough, consider a value of $\theta_\tau \in C_\beta \cap B_{\delta_r}(0)$ with $\|y_\tau\| \neq 0$.

[22] To match the notation of that part of the proof in Evans and Honkapohja (2000), set $b_1 \equiv s - 1 - \delta_a$, $b_2 \equiv a - 1 + \delta_a$, and let $a \equiv b - 1 - \delta_b$, where the right-hand-side variables in these definitions are in our

It must be that for $t \geq \tau$

$$\|\theta_{t+1}\|^2 \geq (1 + \gamma_{t+1} 2 (b - 1 - 2\delta_b)) \|\theta_t\|^2 .$$

Now consider $\tilde{\theta}_t = \frac{\theta_t}{t^{b-1-\delta}}$. For $t \geq \tau$,

$$\left\|\tilde{\theta}_{t+1}\right\|^2 \geq \left(1 + \frac{1}{t}\right)^{2(1-b+\delta)} (1 + \gamma_{t+1} 2 (b - 1 - 2\delta_b)) \left\|\tilde{\theta}_t\right\|^2 .$$

For some finite $C > 0$ and $t \geq \tau$

$$\begin{aligned} \left\|\tilde{\theta}_{t+1}\right\|^2 &> \left(1 + 2 (1 - b + \delta) \frac{1}{t} - C \frac{1}{t^2}\right) (1 + \gamma_{t+1} 2 (b - 1 - 2\delta_b)) \left\|\tilde{\theta}_t\right\|^2 \\ &> (1 + \tilde{\omega}) \left\|\tilde{\theta}_t\right\|^2 \end{aligned}$$

for some $\tilde{\omega} > 0$ to long as τ is large enough. Note that the first inequality follows from Taylor's theorem and the second inequality follows because we chose δ_b so that $2\delta_b < \delta$. So, it must be that

$$\lim_{j \to \infty} \left\|\tilde{\theta}_{\tau+j}\right\| = \infty.$$

The argument involving the inverse function theorem given at the end of the proof of Proposition 2 in Evans and Honkapohja (2000) shows that any open ball around the origin contains a θ_0 that will generate a sequence θ_t so that for $\theta_\tau \in C_\beta \cap B_{\delta_r}(0)$ for some large enough τ, which completes the proof. □

B Variant of CEE (2005) used in this paper

[INCOMPLETE]

C Solution algorithm for non-linear NK model

In this Appendix we detail our solution strategy for the non-linear NK model we consider in our paper. We exploit the model's structure to simplify its solution. In particular, because the steady state is an absorbing state for the REE and the learning equilibria that we consider, we can solve the steady state decision rules without reference to the period when $r = r_\ell$. With this solution in had, we then turn to the period when $r = r_\ell$, which is where we consider learning.

notation. Also, the roles of x_t and y_t are reversed in our notation relative to Evans and Honkapohja (2000).

Our main model code is implemented in c++, with reliance on the Eigen, boost, and nlopt libraries. Our computations were conducted using nearly 400 processors with heavy reliance on MPI. Our computations took roughly two weeks to complete. Details related to our model code are available in the README file associated with the replication materials. This Appendix outlines the strategy used to solve the model that is implemented in that code.

C.1 Steady state

In the steady state, there is no uncertainty. However, households still face a bond-holding choice and firms still face a relative-price choice. In an REE, households will choose to hold zero bonds and firms will choose to set their price to the aggregate price level.

C.1.1 Household problem

In the steady state, the household value function is given by

$$V_{h,ss}\left(b_h\right) = \max_{C_h, N_h, b'_h} \left\{ \log\left(C_h\right) - \frac{\chi}{2}\left(N_h\right)^2 + \beta V_{h,ss}\left(b'_h\right) \right\}$$

subject to

$$C_h + \frac{b'_h}{R_{ss}} \leq \frac{b_h}{\pi_{ss}} + w_{ss} N_h + T_{ss}.$$

Here, b_h and b'_h are household h's real bond holdings chosen in the previous and current period, respectively. The variables C_h and N_h are household h's consumption and labor supply. The aggregate variable R_{ss}, π_{ss}, w_{ss}, and T_{ss} are the gross nominal interest rate, the gross inflation rate, the real wage, and taxes net of transfers and profits. The values of these aggregate variables are known to the household. We constrain households so that $b'_h \in \left[\underline{b}, \overline{b}\right]$. Implicitly, we have functions $C_{h,ss}\left(b_h\right)$, $N_{h,ss}\left(b_h\right)$, and $b'_{h,ss}\left(b_h\right)$. Assuming the constraint on b'_h is not binding, household maximization implies

$$\frac{1}{C_{h,ss}\left(b_h\right)} = \beta R_{ss} \frac{1}{C_{h,ss}\left(b'_h\left(b_h\right)\right)\pi_{ss}} \tag{26}$$

$$\chi C_{h,ss}\left(b_h\right) N_{h,ss}\left(b_h\right) = w_{ss} \tag{27}$$

$$C_{h,ss}\left(b_h\right) + \frac{b'_{h,ss}\left(b_h\right)}{R_{ss}} = \frac{b_h}{\pi_{ss}} + w_{ss} N_{h,ss}\left(b_h\right) + T_{ss} \tag{28}$$

We define a grid over $\left[\underline{b},\overline{b}\right]$ and approximate the functions $C_{h,ss}\left(b_h\right)$, $N_{h,ss}\left(b_h\right)$, and $b'_{h,ss}\left(b_h\right)$ on that grid in the following way.[23]

1. We conjecture a value for $b'_{h,ss}\left(b_h\right)$ at each grid point. Call the conjectured value $b'^{i}_{h,ss}\left(b_h\right)$.

2. Note that equations (27) and (28) can be written as

$$\chi C_{h,ss}\left(b_h\right)\left(C_{h,ss}\left(b_h\right)+\frac{b'^{i}_{h,ss}\left(b_h\right)}{R_{ss}}-\frac{b_h}{\pi_{ss}}-T_{ss}\right)=w_{ss}^2.$$

   The left-hand-side is increasing in $C_{h,ss}\left(b_h\right)\geq 0$. For every b_h, we solve for the value of $C_{h,ss}\left(b_h\right)$ that makes this hold with equality. We call this the conjectured value for $C_{h,ss}\left(b_h\right)$ and denote it by $C^{i}_{h,ss}\left(b_h\right)$. Note that with $C^{i}_{h,ss}\left(b_h\right)$, we can back out $N^{i}_{h,ss}\left(b_h\right)$ from equation (27).

3. For each grid point, b_h, find b'_h that solves the following version of equation (26)

$$C_h\beta R_{ss}\frac{1}{C^{i}_{h,ss}\left(b'_h\right)}-1=0$$

   where $C_h\geq 0$ solves

$$\chi C_h\left(C_h+\frac{b'_h}{R_{ss}}-\frac{b_h}{\pi_{ss}}-T_{ss}\right)=w_{ss}^2.$$

   We use linear interpolation to compute $C^{i}_{h,ss}\left(b'_h\right)$ for values of b'_h that fall between grid points. If the procedure would set $b'_h>\overline{b}$ or $b'_h<\underline{b}$, we set $b'_h$to the respective endpoint of the grid. We record the value of b'_h in by updating the conjectured rule for $b'_{h,ss}\left(b_h\right)$ using $b'^{i+1}_{h,ss}\left(b_h\right)=b'_h$.

4. Having computed $b'^{i+1}_{h,ss}\left(b_h\right)$ for every grid point, we check to see if

$$\left|b'^{i+1}_{h,ss}\left(b_h\right)-b'^{i}_{h,ss}\left(b_h\right)\right|<\epsilon$$

   at every grid point for some small ϵ. If yes, we say that we have solved the household problem in steady state. If no, we set $b'^{i}_{h,ss}\left(b_h\right)=b'^{i+1}_{h,ss}\left(b_h\right)$ and repeat steps (ii), (iii), and (iv).

Because $\beta\frac{R_{ss}}{\pi_{ss}}=1$, it is not surprising that we find that $b'_{h,ss}\left(b_h\right)=b_h$.

[23] In our implementation, we set $-\underline{b}=\overline{b}=1$, which is equal to steady state output. We use a symmetric grid with 25 points that includes zero and places more points near zero than at more extreme values because $b_h=b'_h=0$ in both REE and in learning equilibria.

C.1.2 Firm problem

In the steady state, the firm value function is given by

$$\begin{aligned} V_{f,ss}\left(p_f\right) = \max_{p'_f} \Bigg\{ & \frac{1}{C_{ss}}\left(p'_f - \left(1-\tau_f\right)w_{ss}\right)\left(p'_f\right)^{-\varepsilon} Y_{ss} \\ & - \frac{1}{C_{ss}}\frac{\phi}{2}\left(\frac{p'_f}{p_f}\pi_{ss} - 1\right)^2 \left(C_{ss}+G_{ss}\right) \\ & + \beta V_{f,ss}\left(p'_f\right)\Bigg\}. \end{aligned}$$

Here, p_f and p'_f are the ratio of firm f's price to the aggregate price level in the previous and current period, respectively. The aggregate values π_{ss}, w_{ss}, C_{ss}, G_{ss}, and Y_{ss} are known to the firm. We constrain firms so that $\log\left(p'_f\right) \in \left[\underline{p}, \overline{p}\right]$. Implicitly, we have a function $p'_{f,ss}\left(p_f\right)$. Assuming the constraint on p'_f is not binding, firm maximization implies

$$\begin{aligned} & \phi\left(\frac{p'_f\left(p_f\right)}{p_f}\pi_{ss} - 1\right)\frac{1}{p_f}\pi_{ss}\left(C_{ss}+G_{ss}\right) = \\ & (\varepsilon - 1)\left(\frac{w_{ss}}{p'_f\left(p_f\right)} - 1\right)\left(p'_f\left(p_f\right)\right)^{-\varepsilon} Y_{ss} \\ & + \beta\phi\left(\frac{p'_f\left(p'_f\left(p_f\right)\right)}{p'_f\left(p_f\right)}\pi_{ss} - 1\right)\frac{p'_f\left(p'_f\left(p_f\right)\right)}{\left(p'_f\left(p_f\right)\right)^2}\pi_{ss}\left(C_{ss}+G_{ss}\right) \end{aligned} \tag{29}$$

We define a grid over $\left[\underline{p}, \overline{p}\right]$ and approximate the function $p'_{f,ss}\left(p_f\right)$ on that grid in the following way.[24]

1. We conjecture a value for $p'_{f,ss}\left(p_f\right)$ at each grid point. Call the conjectured value $p'^{i}_{f,ss}\left(p_f\right)$.

[24] In our implementation, we set $-\underline{p} = \overline{p} = 1$. We use a symmetric grid with 25 points that includes zero that places more points near zero than at more extreme values because $\log\left(p_f\right) = \log\left(p'_f\right) = 0$ in both REE and in learning equilibria.

2. For each grid point, p_f, find p'_f that solves the following version of equation (29)

$$\phi\left(\frac{p'_f}{p_f}\pi_{ss}-1\right)\frac{1}{p_f}\pi_{ss}\left(C_{ss}+G_{ss}\right)=$$
$$(\varepsilon-1)\left(\frac{w_{ss}}{p'_f}-1\right)\left(p'_f\right)^{-\varepsilon}Y_{ss}$$
$$+\beta\phi\left(\frac{p'^{i}_{f,ss}\left(p'_f\right)}{p'_f}\pi_{ss}-1\right)\frac{p'^{i}_{f,ss}\left(p'_f\right)}{\left(p'_f\right)^2}\pi_{ss}\left(C_{ss}+G_{ss}\right)$$

We use linear interpolation over $\log\left(p'_f\right)$ to compute $p'^{i}_{f,ss}\left(p'_f\right)$for values of $\log\left(p'_f\right)$ that fall between grid points. If the procedure would set $\log\left(p'_f\right)>\overline{p}$ or $\log\left(p'_f\right)<\underline{p}$, we set $p'_f$to the respective endpoint of the grid. We record the value of p'_f in by updating the conjectured rule for $p'_{f,ss}\left(p_f\right)$ using $p'^{i+1}_{f,ss}\left(p_f\right)=p'_f$.

3. Having computed $p'^{i+1}_{f,ss}\left(p_f\right)$ for every grid point, we check to see if

$$\left|p'^{i+1}_{f,ss}\left(p_f\right)-p'^{i}_{f,ss}\left(p_f\right)\right|<\epsilon$$

at every grid point for some small ϵ. If yes, we say that we have solved the firm problem in steady state. If no, we set $p'^{i}_{f,ss}\left(p_f\right)=p'^{i+1}_{f,ss}\left(p_f\right)$ and repeat steps (ii) and (iii).

C.2 Solution when $r=r_\ell$

To address the case when $r=r_\ell$, we assume that we have the steady state decision rules in hand and that households and firms know these decision rules with certainty.

C.2.1 Beliefs

Before presenting the household and firm problems, some comments about beliefs are in order when $r=r_\ell$. To simplify the model, we assume households and firms have the same beliefs (though they do not know that they have the same beliefs). Households and firms believe that so long as $r=r_\ell$ the log of aggregate consumption, $\log(C)$, and the log of aggregate inflation, $\log(\pi)$, have uncorrelated Normal distributions with unknown means and variances. That is

$$\log(\pi)\sim N\left(\mu_\pi,\sigma_\pi^2\right)$$
$$\log(C)\sim N\left(\mu_C,\sigma_C^2\right).$$

We assume that households and firms have beliefs about the means and variances of the distributions for $\log(C)$ and $\log(\pi)$ that are characterized by density functions that are proportional to Normal-inverse-gamma distributions. These beliefs are not exactly Normal-inverse-gamma distributions because the households and firms embed in their beliefs an upper bound on the variances. This upper bound is important because if variances were unbounded, $\mathbb{E}[\pi] = \mathbb{E}[C] = \infty$, which would challenge the applicability of an expected utility framework. The distributions characterizing beliefs are independent across C and π. That is, for $i \in \{\pi, C\}$, $\mu_i \in (-\infty, \infty)$ and $\sigma_i^2 \in [0, \overline{\sigma}_i^2]$ we have

$$\Pr\left(\sigma_i^2 | \alpha_i, \beta_i\right) = \frac{\frac{\beta_i^{\alpha_i}}{\Gamma(\alpha_i)} \left(\frac{1}{\sigma_i^2}\right)^{\alpha_i+1} \exp\left(-\frac{\beta_i}{\sigma_i^2}\right)}{\frac{\Gamma\left(\alpha_i, \beta_i/\overline{\sigma}_i^2\right)}{\Gamma(\alpha_i)}},$$
$$\Pr\left(\mu_i | \sigma_i^2, m_i, \lambda_i\right) = \frac{\sqrt{\lambda_i}}{\sqrt{2\pi\sigma_i^2}} \exp\left(-\frac{\lambda_i}{2\sigma_i^2}\left(\mu_i - m_i\right)^2\right).$$

Here, $\Gamma(\cdot)$ is the gamma function and $\Gamma(\cdot,\cdot)$ is the incomplete gamma function. Note that $\Gamma(\cdot) = \Gamma(\cdot, 0)$. Again, the advantage of truncating the support of σ_i^2 is that $\mathbb{E}[\pi] < \infty$ if $\overline{\sigma}_\pi^2 < \infty$ and $\mathbb{E}[C] < \infty$ if $\overline{\sigma}_C^2 < \infty$.

Even though we truncate the distributions for σ_i^2, we maintain conjugacy between prior and posterior beliefs as well as the usual recursive updating equations because the likelihoods associated with observations of π and C are not truncated. Beliefs about $\log(i)$ are parameterized by four values, α_i, β_i, m_i, and λ_i. So, we have 8 total values for both π and C. The standard recursive updating formulas for these variables are

$$\lambda_i' = \lambda_i + 1$$
$$m_i' = \frac{\lambda m_i + \log(i)}{\lambda + 1}$$
$$\alpha_i' = \alpha_i + 1/2$$
$$\beta_i' = \beta_i + \frac{\lambda_i\left(\log(i) - m_i\right)^2}{2\left(\lambda_i + 1\right)}.$$

Here, a prime indicates the value taken after having observed $\log(i)$.

We need to include variables in Θ that will fully capture the values α_i, β_i, m_i, and λ_i for $i \in \{\pi, C\}$. First, we keep $\frac{1}{t_\ell}$ in Θ, which is the inverse of the number of periods that r has been equal to r_ℓ. We keep the inverse because it is bounded between zero and one, which will be useful. From this value, we can trivially back out λ_i and α_i, given their values in the first period when $r = r_\ell$. We set the initial value of $\lambda_i = 1$ and the initial value of $\alpha_i = 2$.

We keep m_C and m_π in Θ. And we also keep

$$\psi_i' = \sqrt{\psi_i^2 \frac{2\alpha_i'}{2\alpha_i'+1} + \frac{\lambda_i}{\lambda_i+1}\frac{1}{2\alpha_i'+1}\left(\log(i) - m_i\right)^2}.$$

Note that by setting $\beta_i = (\psi_i)^2 \alpha_i'$ it is clear that we recover the exactly recursive structure of β_i (given above). An advantage of using ψ_i in Θ rather than β_i is that ψ_i is a consistent estimator for the standard deviation, whereas β_i generally grows without bound (except when the standard deviation is zero). Keeping the values of Θ within bounded grids will be important for the purposes of approximation. In total, $\Theta = \left[\frac{1}{\tau_\ell}, m_\pi, m_C, \psi_\pi, \psi_C\right]$ has five elements and we have a mapping from Θ to α_i, β_i, m_i, and λ_i for $i \in \{\pi, C\}$. We also have a law of motion for Θ so that $\Theta' = L(\Theta, [\pi, C])$.

An advantage of the Normal-inverse-gamma setup detailed above is that we can have analytic expressions for the distribution for the variables $\log(\pi)$ and $\log(C)$ conditional on Θ. In particular

$$\begin{aligned}
\Pr(\log(i)|\Theta) =& \frac{\Pr(\log(i)|\mu_i, \sigma_i^2, \Theta)\Pr(\mu_i, \sigma_i^2|\Theta)}{\Pr(\mu_i, \sigma_i^2|\log(i), \Theta)} \\
=& \frac{\frac{1}{\sqrt{2\pi\sigma_i^2}}\exp\left(-\frac{1}{2\sigma_i^2}\left(\log(i) - m_i\right)^2\right)}{\frac{\sqrt{\lambda_i'}}{\sqrt{2\pi\sigma_i^2}}\exp\left(-\frac{\lambda_1'}{2\sigma_i^2}\left(\mu_i - m_i'\right)^2\right)\frac{(\beta_i')^{\alpha_i'}}{\Gamma(\alpha_i')}\left(\frac{1}{\sigma_i^2}\right)^{\alpha_i'+1}\exp\left(-\frac{\beta_i'}{\sigma_i^2}\right)} \\
&\times \frac{\sqrt{\lambda_i}}{\sqrt{2\pi\sigma_i^2}}\exp\left(-\frac{\lambda_i}{2\sigma_i^2}\left(\mu_i - m_i\right)^2\right) \\
&\times \frac{\beta_i^{\alpha_i}}{\Gamma(\alpha_i)}\left(\frac{1}{\sigma_i^2}\right)^{\alpha_i+1}\exp\left(-\frac{\beta_i}{\sigma_i^2}\right)\frac{\kappa_i'}{\kappa_i}
\end{aligned}$$

where

$$\kappa_i = \frac{\Gamma\left(\alpha_i, \beta_i/\overline{\sigma}^2\right)}{\Gamma(\alpha_i)}.$$

Then

$$\begin{aligned}
\Pr(\log(i)|\Theta) =& \left(\frac{\lambda_i\alpha_i}{\beta_i(\lambda_i+1)}\right)^{1/2}\frac{\Gamma(\alpha_i + 1/2)}{\Gamma(\alpha_i)\sqrt{2\pi\alpha_i}} \\
&\times \left(1 + \frac{1}{2\alpha_i}\frac{\left(\log(i) - m_i\right)^2}{\left(\frac{\lambda_i\alpha_i}{\beta_i(\lambda_i+1)}\right)^{-1}}\right)^{-\alpha_0 - 1/2}\frac{\kappa_i'}{\kappa_i}. \qquad (30)
\end{aligned}$$

Notice that κ_i' depends on the point of evaluation for $\log(i)$. Evidently, if we ignored the ratio κ_i'/κ_i, which would be correct in the case when $\bar{\sigma}_i^2 = \infty$, the pdf for $\log(i)$ is a t distribution with location parameter m_i, scale parameter $\left(\frac{\lambda_i \alpha_i}{\beta_i(\lambda_i+1)}\right)^{-1/2}$, and $2\alpha_i$ degrees of freedom. If $\overline{\sigma}_i^2$ is large, $\kappa_i'/\kappa_i \neq 1$ but is close to unity. For finite $\bar{\sigma}_i^2$, the ratio κ_i'/κ_i serves to thin the tails of the distribution of $\log(i)$ by down-weighting the probability of extreme values for $\log(i)$.[25] Because the density function of the t distribution is readily available and reliably computed in statistical software and because κ_i and κ_i' are easily computed using readily available implementations of the gamma and incomplete gamma functions, we can use equation (30) for quadrature weighting. We use Gauss-Hermite quadrature with seven nodes when computing approximations to integrals based on equation (30).

C.2.2 Household problem

When $r = r_\ell$, the household value function is given by

$$V_h(b_h, \Theta, x) = \max_{C_h, N_h, b_h'} \left\{ \log(C_h) - \frac{\chi}{2}(N_h)^2 + \frac{1}{1+r_\ell}\left[p\mathbb{E}_{\Theta'} V_h(b_h', \Theta', x') + (1-p)V_{h,ss}(b_h')\right] \right\}$$

subject to

$$C_h + \frac{b_h'}{R} \leq \frac{b_h}{\pi} + wN_h + T.$$

Here, $x = [\pi, C]'$, $V_{h,ss}(\cdot)$ is the steady state value function for the household, which is defined above, and $\mathbb{E}_{\Theta'}$ denotes expectations of the household computed conditional on Θ'. Given x and Θ, we have $\Theta' = L(\Theta, x)$. So, the expectation of the household is taken with respect to x', which is believed to be iid. We assume that households know the monetary and fiscal policy rules. We also assume that they correctly think that $Y = (C+G)\left(1 + \frac{\phi}{2}(\pi-1)^2\right)$, $N = Y$, and $w = \chi CY$. Given x, with these assumptions R, π, w, and T can be computed. The steady state values of aggregate variables are known to the household. We constrain households so that $b_h' \in \left[\underline{b}, \overline{b}\right]$. The household optimization problem gives us implicit functions for $C_h(b_h, \Theta, x)$, $N_h(b_h, \Theta, x)$, and $b_h'(b_h, \Theta, x)$. Considering interior solutions for b_h',

[25] We set $\bar{\sigma}_i^2$ equal to the squared maximum value on the grid for ψ_i (described below).

we have

$$\frac{1}{C_h(b_h,\Theta,x)} = \frac{1}{1+r_\ell}R\left[p\mathbb{E}_{\Theta'}\left\{\frac{1}{\pi' C_h'(b_h',\Theta',x')}\right\} + (1-p)\frac{1}{\pi_{ss}C_{h,ss}(b_h')}\right] \tag{31}$$

$$w = \chi C_h(b_h,\Theta,x)\,N_h(b_h,\Theta,x) \tag{32}$$

$$C_h(b_h,\Theta,x) = \frac{b_h}{\pi} + wN_h(b_h,\Theta,x) + T - \frac{b_h'(b_h,\Theta,x)}{R}. \tag{33}$$

Instead of approximating $C_h(b_h,\Theta,x)$, $N_h(b_h,\Theta,x)$, and $b_h'(b_h,\Theta,x)$ directly, we approximate

$$v_h(b_h,\Theta) = \mathbb{E}_\Theta\left\{\frac{1}{\pi C_h(b_h,\Theta,x)}\right\}.$$

We take this approach because we can eliminate x as a state variable in the approximation. We define grids on the elements of Θ and use the grid defined for b_h in the steady state. We then approximate $v_h(b_h,\Theta)$ in the following way.[26]

1. We conjecture a value for $v_h(b_h,\Theta)$ at each grid point in the cross product of the grids over the elements of b_h and Θ.[27] Call the conjectured value $v_h^i(b_h,\Theta)$.

2. For a given grid point we use quadrature to get a value for $\mathbb{E}_\Theta\left\{(\pi C_h)^{-1}\right\}$.To solve for the expectation of interest, we need to solve for C_h given many different values for x. Conditional on a value for x, equations (32) and (33) can be written as

$$\chi C_h\left(C_h + \frac{b_h'}{R} - \frac{b_h}{\pi} - T\right) = w^2$$

The left-hand-side is increasing in $C_h \geq 0$. For a given b_h', we solve for the value of C_h that makes this hold with equality. We then search for the value of b_h' that makes the following version of equation (31) hold with equality:

$$\frac{1}{C_h} = \frac{1}{1+r_\ell}R\left[pv_h^i(b_h',\Theta') + (1-p)\frac{1}{\pi_{ss}C_{h,ss}(b_h')}\right].$$

We use linear interpolation to compute $v_h^i(b_h',\Theta')$ for values of b_h' and Θ' that fall

[26]The grids for m_i contain 12 points that are are not evenly spaced. They include each REE point as well as the target-inflation steady state. The remaning points are bunched relatively close to the REE points. The grid for ψ_C contains 11 points that are evenly spaced from 0 to 0.1. The grid for ψ_π contains 11 points that are evenly spaced from 0 to 0.05. Note that inflation is expressed in quarterly terms, so a change of 0.05 would be 20 percent if annualized.

[27]There are $435,600 = 12\times12\times11\times11\times25$ points in the cross product of the grids for m_i, ψ_i, and b_h. The grid for t_ℓ^{-1} is handled in a way discussed below.

between grid points. If the procedure would set $b'_h > \bar{b}$ or $b'_h < \underline{b}$, we set $b'_h$to the respective endpoint of the grid for b_h. We record the associated value of C_h and use it in the quadrature to compute $v_h^{i+1}(b_h, \Theta) = \mathbb{E}_\Theta \left\{ (\pi C_h)^{-1} \right\}$.

3. Having computed $v_h^{i+1}(b_h, \Theta)$ for every grid point, we check to see if

$$\left| v_h^i (b_h, \Theta) - v_h^{i+1} (b_h, \Theta) \right| < \epsilon$$

at every grid point for some small ϵ. If yes, we say that we have solved the household problem when $r = r_\ell$. If no, we set $v_h^i(b_h, \Theta) = v_h^{i+1}(b_h, \Theta)$, repeat steps (ii) and (iii).

The grid that we use on $\frac{1}{t_\ell}$ is special. In particular, we let that grid be $\left[0, \frac{1}{99}, \frac{1}{98}, \ldots, 1\right]$. The first element of the grid corresponds to the case when infinite time has past. In this case households think that they would update their beliefs so that $\Theta' = \Theta$ because m_i and ψ_i are consistent estimators for the means and variances. In our numerical computations, we utilize this fact to first approximate v_h in this case. We then approximate v_h in the case where $t_\ell = 99$. When $t_\ell = 99$, we need to interpolate between the solution to the case when $t_\ell = \infty$ and the conjectured value of $v_h^i(b_h, \Theta)$ when $t_\ell = 99$ to evaluate $v_h^i(b'_h, \Theta')$. That is, when $t_\ell = 99$ we have to find a fixed point of this interpolation, which is computationally intense. To do the interpolation, we linearly interpolate between $t_\ell^{-1} = 1/99$ and $t_\ell^{-1} = 0$. When $t_\ell = 98$, having approximated $v_h(b_h, \Theta)$ for $t_\ell = 99$ means that can evaluate $v_h(b'_h, \Theta')$ exactly at $t'_\ell = 99$ without reference to $v_h^i(b_h, \Theta)$. We approximate for $v_h(b_h, \Theta)$ when $t_\ell = 98$ and work work back in this way to $t_\ell = 1$. This strategy fits this into the structure of steps 1-3 because we know that the value of $v_h(b_h, \Theta)$ will not depend on its value at any any t_ℓ that is smaller than implied by Θ. So, we have a block dependent structure to $v_h(b_h, \Theta)$. Additionally, we know that t_ℓ will only take integer values.

C.2.3 Firm problem

When $r = r_\ell$, the firm value function is given by

$$V_f(p_f, \Theta, x) = \max_{p'_f} \left\{ \frac{1}{C} \left(\left(p'_f - (1 - \tau_f) w \right) \left(p'_f \right)^{-\varepsilon} Y - \frac{\phi}{2} \left(\frac{p'_f}{p_f} \pi - 1 \right)^2 (C + G) \right) \right.$$
$$\left. + \frac{1}{1 + r_\ell} \left[p \mathbb{E}_{\Theta'} V_f \left(p'_f, \Theta', x' \right) + (1 - p) V_{f,ss} \left(p'_f \right) \right] \right\}$$

Here, $x = [\pi, C]'$, $V_{f,ss}(\cdot)$ is the steady state value function for the firm, which is defined above, and $\mathbb{E}$ denotes expectations of the firm. Given x and Θ, we have $\Theta' = L(\Theta, x)$.

So, the expectation of the firm is taken with respect to x', which is believed to be iid. We assume that firms know the monetary and fiscal policy rules. We also assume that they correctly think that $Y = (C+G)\left(1+\frac{\phi}{2}(\pi-1)^2\right)$, $N = Y$, and $w = \chi CY$. Given x, with these assumptions π, w, G, and Y can be computed. The steady state values of aggregate variables are known to the firm. We constrain firms so that $\log\left(p'_f\right) \in \left[\underline{p}, \overline{p}\right]$. Implicitly, from firm optimization we have a function $p'_f(p_f, \Theta, x)$. Considering interior solutions for p'_f, firm maximization implies

$$
\begin{aligned}
&\phi\left(\frac{p'_f(p_f,\Theta,x)}{p_f}\pi-1\right)\frac{1}{p_f}\pi(C+G)= \\
&(\varepsilon-1)\left(\frac{w}{p'_f(p_f,\Theta,x)}-1\right)\left(p'_f(p_f,\Theta,x)\right)^{-\varepsilon}Y \\
&+\frac{1}{1+r_\ell}p\mathbb{E}_{\Theta'}\frac{C}{C'}\phi\left(\frac{p'_f\left(p'_f(p_f,\Theta,x),\Theta',x'\right)}{p'_f(p_f,\Theta,x)}\pi'-1\right)\frac{p'_f\left(p'_f(p_f,\Theta,x),\Theta',x'\right)}{\left(p'_f(p_f,\Theta,x)\right)^2}\pi'(C'+G') \\
&+\frac{1}{1+r_\ell}\frac{C}{C_{ss}}(1-p)\phi\left(\frac{p'_{f,ss}\left(p'_f(p_f,\Theta,x)\right)}{p'_f(p_f,\Theta,x)}\pi_{ss}-1\right)\frac{p'_{f,ss}\left(p'_f(p_f,\Theta,x)\right)}{\left(p'_f(p_f,\Theta,x)\right)^2}\pi_{ss}(C_{ss}+G_{ss}).
\end{aligned}
\tag{34}
$$

Instead of approximating $p'_f(p_f,\Theta,x)$ directly, we approximate

$$v_f(p_f,\Theta) = \mathbb{E}_\Theta\left\{\frac{1}{C}\phi\left(\frac{p'_f}{p_f}\pi-1\right)\frac{p'_f}{p_f}\pi(C+G)\right\}.$$

We take this approach because we can eliminate x as a state variable in the approximation. We use the same grids on the elements of Θ that we use for the household problem and the grid defined for $\log(p_f)$ in the steady state and we approximate $v_f(p_f,\Theta)$ in the following way.

1. We conjecture a value for $v_f(p_f,\Theta)$ at each grid point in the cross product of the grids over the elements of p_f and Θ. Call the conjectured value $v^i_f(p_f,\Theta)$.

2. For a given grid point we use quadrature to get a value for

$$\mathbb{E}\left\{\frac{1}{C}\phi\left(\frac{p'_f}{p_f}\pi-1\right)\frac{p'_f}{p_f}\pi(C+G)\right\}.$$

   To solve for the expectation of interest, we need to solve for p'_f given many different values for x. Conditional on a value for x, we find a value of p'_fthat solves the following version of equation (34)

$$\begin{aligned}
&\phi\left(\frac{p_f'}{p_f}\pi-1\right)\frac{1}{p_f}\pi\left(C+G\right)= \\
&\left(\varepsilon-1\right)\left(\frac{w}{p_f'}-1\right)\left(p_f'\right)^{-\varepsilon}Y \\
&+\frac{1}{1+r_\ell}pv_f^i\left(p_f',\Theta'\right)\frac{C}{p_f'} \\
&+\frac{1}{1+r_\ell}\frac{C}{C_{ss}}\left(1-p\right)\phi\left(\frac{p_{f,ss}'\left(p_f'\right)}{p_f'}\pi_{ss}-1\right)\frac{p_{f,ss}'\left(p_f'\right)}{\left(p_f'\right)^2}\pi_{ss}\left(C_{ss}+G_{ss}\right).
\end{aligned}$$

We use linear interpolation over $\log\left(p_f'\right)$ to compute $v_f^i\left(p_f',\Theta'\right)$ for values of $\log\left(p_f'\right)$ and Θ' that fall between grid points. If the procedure would set $\log\left(p_f'\right)>\overline{p}$ or $\log\left(p_f'\right)<\underline{p}$, we set $p_f'$to the respective endpoint of the grid for p_f. We record the value of p_f' in and the associated aggregate variables so that the quadrature procedure can approximate

$$v_f^{i+1}\left(p_f,\Theta\right)=\mathbb{E}_\Theta\left\{\frac{1}{C}\phi\left(\frac{p_f'}{p_f}\pi-1\right)\frac{p_f'}{p_f}\pi\left(C+G\right)\right\}.$$

3. Having computed $v_f^{i+1}\left(p_f,\Theta\right)$ for every grid point, we check to see if

$$\left|v_f^i\left(p_f,\Theta\right)-v_f^{i+1}\left(p_f,\Theta\right)\right|<\epsilon$$

at every grid point for some small ϵ. If yes, we say that we have solved the household problem when $r=r_\ell$. If no, we set $v_f^i\left(p_f,\Theta\right)=v_f^{i+1}\left(p_f,\Theta\right)$, repeat steps (ii) and (iii).

Our use of the same grids as in the household problem allows us to exploit the same block dependent structure in t_ℓ^{-1}.

C.3 Learning equilibria

Here we detail how we construct learning equilibria, given the solutions to the household an firm problems—v_h and v_f.

1. Set $r=r_\ell$ and assume a value for Θ_t for $t=1$.

2. Conjecture a value for π_t.

(a) Find the value of C_t that would make the following equation hold

$$\frac{1}{C_t} = \frac{1}{1+r_\ell} R_t \left[pv_h \left(0, f\left(\Theta_t, [\pi_t, C_t]\right)\right) + (1-p) \frac{1}{\pi_{ss} C_{h,ss}(0)} \right].$$

Note that with π_1 and C_1 the values of all other aggregate variables can be computed.

(b) Check to see if the following equation holds

$$\begin{aligned}
&\phi(\pi_t - 1)\pi_t (C_t + G_t) = \\
&\quad (\varepsilon - 1)(w_t - 1) + \frac{1}{1+r_\ell} pv_f \left(1, f\left(\Theta_t, [\pi_t, C_t]\right)\right) C_t \\
&\quad + \frac{1}{1+r_\ell} \frac{C_t}{C_{ss}} (1-p) \phi(\pi_{ss} - 1) \pi_{ss} (C_{ss} + G_{ss}).
\end{aligned}$$

If yes, we have a period equilibrium for period t and we record π_t and C_t. If no, conjecture a different value for π_t.

3. Set $\Theta_{t+1} = L\left(\Theta_t, [\pi_t, C_t]\right)$ and repeat step (ii).

When we consider "anticipated utility," we define $\tilde{\Theta}_t$ to be Θ_t, but with $\frac{1}{t_\ell} = 0$. We then perform step 2 with $\tilde{\Theta}_t$ instead of Θ_t. However, in step 3 we continue to use Θ_t. The switch between $\tilde{\Theta}_t$ and Θ_t highlights the way in which "anticipated utility" is not internally rational.

D Linearization under anticipated utility and learning about intercepts

Here we describe our strategy for linearizing models under anticipated utility and learning. This strategy applies to our simple NK model without the ZLB and to our variant of the CEE (2005) model. To fix notation, if a variable, x_t has a steady state value x that is different from zero, then we define $\widehat{x}_t = \log(x_t/x)$. Otherwise, we define $\widehat{x}_t = x_t$.

We assume that there is a unique, bounded linear REE in which

$$\widehat{a}_t = \Theta_{1,RE}\widehat{a}_{t-1} + \Theta_{\epsilon,RE}\epsilon_t. \tag{35}$$

Here, $\widehat{a}_t$ is a vector of aggregate variables in (log-)deviation from steady state and ϵ_t is a (possibly empty) vector of standard Normal structural shocks. In equation 35, and throughout this appendix, we use notation similar to that used in Sims (2001). All matrices are assumed to be conformable.

In our learning model, agents have beliefs about a vector of aggregate variables, $x_t \subseteq a_t$. We assume that those beliefs are of the form

$$\widehat{x}_t = c_x + \Gamma_{1,x}\widehat{a}_{t-1} + \Psi_{x,\epsilon}\epsilon_t + \Psi_{x,\nu}\nu_t. \tag{36}$$

Agents are uncertain about c_x and $\Psi_{x,\nu}$. We assume that agents know $\Gamma_{1,x}$ and $\Psi_{x,\epsilon}$ to be their REE values, consistent with equation 35. Agents also think that there is another vector of shocks, ν_t, that could affect $\widehat{x}$. Agents believe that ν_t is a vector of independent standard Normal random variables. We assume that the matrix $\Psi_{x,\nu}$ is diagonal, reflecting the belief that the elements of ν_t only affect the associated element of $\widehat{x}_t$. In an REE each element c_x and $\Psi_{x,\nu}$ is zero. Note that the elements of ν_t do not appear in the dynamics of $\widehat{a}_t$ in equation 35.

Under the anticipated utility assumption, agents believe with certainty that

$$\widehat{x}_t = c_{x,t-1} + \Gamma_{1,x}\widehat{a}_{t-1} + \Psi_{x,\epsilon}\epsilon_t + \Psi_{x,\nu,t-1}\nu_t. \tag{37}$$

Additionally, we assume that agents know static relationships that allow them to uniquely determine a_t given x_t through the relationship $a_t = F(x_t, \epsilon_t, a_{t-1})$. The (log-)linear version of this mapping is given by

$$\widehat{a}_t = \widehat{F}\begin{bmatrix} \widehat{x}_t \\ \epsilon_t \\ \widehat{a}_{t-1} \end{bmatrix}, \tag{38}$$

where $\hat{F}$ is a conformable matrix.

In the model agents make a number of choices conditional on aggregate states and shocks. For example, in the simple NK model, households choose household-specific bond holdings that they contemplate being different from zero. Also, firms choose firm-specific prices that they contemplate as being different from the aggregate price level. We gather the agent-specific variables that they choose into the vector h_t. We specify h_t so that it includes agent-specific versions of each element of x_t. We can do this because we assume that all agents are symmetric in the groups associated with the representative agent model. Notably, every value of h_t either corresponds to an aggregate quantity in a_t or is a constant value in all states in an REE. An example of an element of h_t corresponding to an aggregate quantity in a_t in the simple NK model is household-specific consumption (in h_t) corresponding to aggregate consumption (in a_t). An example of an element of h_t corresponding to a constant value in all state of an REE in the simple NK model is household-specific bond holdings (in

h_t) corresponding to aggregate bonds being in zero net supply.

After (log-)linearization, agent-specific optimality conditions and budget constraints can be written as

$$\Gamma_{0,h}\widehat{h}_t + \Gamma_{0,e}\widehat{e}_t + \Gamma_{0,a}\widehat{a}_t = \Gamma_{1,h}\widehat{h}_{t-1} + \Gamma_{1,ha}\widehat{a}_{t-1} + \Gamma_{1,h\epsilon}\epsilon_t. \tag{39}$$

Here, the vector $\widehat{e}_t = [\mathbb{E}_t\widehat{h}'_{t+1}, \mathbb{E}_t\widehat{a}'_{t+1}]'$ and $\mathbb{E}_t$ is the expectation operator conditional on agents' beliefs. Those beliefs are given by equation 37. Note that we assume that the model does not contain any agent-specific shocks. We emphasize emphasize that all of the matrices in equation 39 take the values they would in an REE. Deviations from rational expectations only appear through the coefficients in equation 37. Defining $\eta_t = [(\widehat{h}_t - \mathbb{E}_{t-1}\widehat{h}_t)', (\widehat{a}_t - \mathbb{E}_{t-1}\widehat{a}_t)']'$, we also need to impose the identity

$$\begin{bmatrix} h_t \\ a_t \end{bmatrix} = e_{t-1} + \eta_t \tag{40}$$

Letting $y_t = [h'_t, a'_t, e'_t]'$ and equations 37-40 can be written as

$$\Gamma_0 y_t = \Gamma_1 y_{t-1} + C_{t-1} + \Psi_\epsilon \epsilon_t + \Psi_{\nu,t-1}\nu_t + \Pi\eta_t. \tag{41}$$

Letting $z_t = [\epsilon'_t, \nu'_t]'$, this is exactly the form of the system studied by Sims (2001). The solution to this system described in Sims (2001) ensures that $\mathbb{E}_t\eta_{t+1} = 0$. The analysis in Sims (2001) is focused on RE models. We put subscripts on C_{t-1} and $\Psi_{\nu,t-1}$ to emphasize that the coefficients in these matrices depend on beliefs that may be different from RE beliefs. It is worth stressing that the methodology we outline here could be used to compute decision rules for h_t in an REE by setting the matrices in equation 37 to be their REE values. However, in our analysis, households and firms do not necessarily have RE. The deviation from RE in their beliefs is captured by $c_{x,t-1}$ and $\Psi_{x,\nu,t-1}$ through C_{t-1} and $\Psi_{\nu,t-1}$.

Under the conditions discussed in Sims (2001), there is a unique solution to the system of equations in 41 that can be expressed as

$$\widehat{y}_t = \Theta_1\widehat{y}_{t-1} + \Theta_C C_{t-1} + \Theta_\epsilon\epsilon_t + \Theta_{\nu,t-1}\nu_t. \tag{42}$$

Importantly, this system contains the mapping

$$\widehat{h}_t = \Theta_{1,h}\widehat{y}_{t-1} + \Theta_{C,h}C_{t-1} + \Theta_{h,\epsilon}\epsilon_t + \Theta_{h,\nu,t-1}\nu_t \tag{43}$$

that defines decision rules for $\widehat{h}_t$ that are optimal from the perspective of the agents in the model, given $\widehat{y}_{t-1}$, ϵ_t, and z_t. It follows from equation 43-45 in Sims (2001) that because Γ_0, Γ_1, Ψ_ϵ, and Π in equation 41, are equal to their values used when computing an REE, $\Theta_{1,h}$ is also equal to its value in an REE. Additionally, it follows that $\Theta_{h,\epsilon}$ is also equal to its value in an REE. So, we do not include a time subscript on $\Theta_{h,\epsilon}$.[28] Unlike in an REE, the system contains perceived shocks, ν_t. The values of ν_t are determined by the requirement that markets clear in the period equilibrium, which we discuss below.

In a learning equilibrium, we first note that the symmetry of the agents' problems implies that any agent-specific values up to period $t-1$ are equal to their aggregate quantities. So, we can write equation 43 as

$$\widehat{h}_t = \Theta_{1,ha}\widehat{a}_{t-1} + \Theta_{C,h}C_{t-1} + \Theta_{h,\epsilon}\epsilon_t + \Theta_{h,\nu,t-1}\nu_t. \tag{44}$$

The vector $\widehat{e}_{t-1}$ does not appear on the right-hand side of equation 44 because in the solution to the linear model its terms are are a linear combination of ν_t and ϵ_t (see Sims (2001)).

To construct a period equilibrium, we impose the restrictions of market clearing. Those restrictions take the form that the elements of $\widehat{x}_t$ must be equal to corresponding elements of $\widehat{h}_t$. An example of such a restriction in the simple NK model is that household-specific consumption must equal aggregate consumption. The number of such restrictions is equal the length of ν_t in order to determine the values of ν_t. We write these restrictions as

$$R\widehat{h}_t = \widehat{x}_t. \tag{45}$$

Here, R is a matrix of ones and zeros that selects the element of $\widehat{h}_t$ that corresponds to the associated element to $\widehat{x}_t$. In a period equilibrium, equation 45 will hold. From the perspective of the agents, we can use equations 37 and 44 so that equation 45 takes the form

$$(R\Theta_{h,\nu,t-1} - \Psi_{x,\nu,t-1})\,\nu_t = c_{x,t-1} - R\Theta_{C,h}C_{t-1}.$$

If $(R\Theta_{h,\nu,t-1} - \Psi_{x,\nu,t-1})$ is invertable, then the perceived values of ν_t can be determined uniquely. These are the values that ν_t takes in the period equilibrium to induce agents to make decisions that lead to market clearing.

In a learning equilibrium, after agents observe $\widehat{a}_t$ and ϵ_t and make their decisions in the period equilibrium, they update their beliefs using Bayes rule. Remember that agents are

[28] To see this, note that Φ in equation 42 in Sims (2001) is determined only by Π, Γ_0, and Γ_1, which are equal to their values used to compute an REE. It follow from equation 45 of Sims (2001) that Θ_1 in equation 42 is equal to its REE value. Additionally, if we define $\Psi = [\Psi_\epsilon, \Psi_{\nu,t-1}]$, then we can use equation 45 of Sims (2001) (and the notation therein) to set $\Theta_\epsilon = \Omega_{22}^{-1}Q_{2\cdot}\Psi_\epsilon$ and $\Theta_{\nu,t-1} = \Omega_{22}^{-1}Q_{2\cdot}\Psi_{\nu,t-1}$. These expressions delivers the claim in the text regarding $\Theta_{h,\epsilon}$.

only learning about c_x and $\Psi_{x,\nu}$ in equation 36 and that the elements of ν_t are perceived to be independent of one another and over time. We assume that agents have the Normal-inverse-gamma conjugate priors on each element of c_x and the associated element of $\Psi_{x,\nu}$ and that they have the same degrees of freedom in these priors.[29] In this case, agents' posterior beliefs about c_x and $\Psi_{x,\nu}$ are Normal-inverse-gamma and the mean of the beliefs about c_x evolves according to

$$c_{x,t} = c_{x,t-1} + \gamma_t \left(\widehat{x}_t - \Gamma_1 \widehat{a}_{t-1} - \Psi_{x,\epsilon}\epsilon_t - c_{x,t-1}\right) = c_{x,t-1} + \gamma_t \Psi_{x,\nu}\nu_t.$$

Here, $\gamma_t = 1/(\lambda_0 + t)$ where λ_0 determines the precision of the prior on c_x. The values of $\Psi_{x,\nu}\nu_t$ are determined by $c_{x,t-1}$, $\Gamma_{1,x}$, and $\Psi_{x,\epsilon}$, but not by $\widehat{a}_{t-1}$, ϵ_t, or $\Psi_{x,\nu,t-1}$. Therefore, in a learning equilibrium in which $\Gamma_{1,x}$ and $\Psi_{x,\epsilon}$ are known with certainty to be their REE values, the mean of agents' beliefs about c_x volves in a non-stochastic manner and $c_{x,t}$ does not depend on $\widehat{a}_{t-1}$. That is, we can write the evolution of $c_{x,t}$ as

$$c_{x,t} = c_{x,t-1} + \gamma_t \left(\mu(c_{x,t-1}) - c_{x,t-1}\right). \tag{46}$$

Equation 46 is of the form of equation 1 in the main text of our paper.

D.1 The period equilibrium in the simple NK model with only sticky prices

In this subsection, we solve for the period equilibrium of the simple NK model with only sticky prices. To do this, we put a (log-)linear approximation to the solution the simple NK model in the for discussed above. In an REE, this model has no aggregate state variables. We display the equations in their (log-)linear form around a steady state. The model has two such steady states (one with target inflation and one with deflation). Our approach can be used to linearize around either and we use general versions of the equations below.

We assume that households and firms believe with certainty that

$$\widehat{C}_t = m_{C,t-1} + \sigma_{\nu,C,t-1}\nu_{C,t} \tag{47}$$

$$\widehat{\pi}_t = m_{\pi,t-1} + \sigma_{\nu,\pi,t-1}\nu_{\pi,t} \tag{48}$$

Because the REE has no aggregate state variables, these beliefs nest the REE beliefs. Define $x_t = [C_t, \pi_t]$ and note that equations 47-48 take the form of those in the vector eqation 37.

[29] At the expense of additional notation, the assumption about the same degrees of freedom across the priors can be relaxed if μ as defined in equation 46 takes γ_t as an argument, which is permitted by equation 1 in the main text. This change has no effect on the asymptotic rate of convergence of $c_{x,t}$.

Households and firms know the following static relationships

$$Y_t = N_t \tag{49}$$

$$Y_t = (C_t + G_t)\left(1 + \frac{\Phi}{2}(\pi_t - 1)^2\right) \tag{50}$$

$$w_t = \chi N_t C_t \tag{51}$$

$$\tau_t = G_t - (1 - w_t)Y_t + \frac{\Phi}{2}(\pi_t - 1)^2(C_t + G_t) \tag{52}$$

$$G_t = G \tag{53}$$

$$R_t = \max\left\{1, \beta^{-1} + \alpha_\pi(\pi_t - 1) + \alpha_Y \frac{Y_t - Y}{Y}\right\}. \tag{54}$$

The first equation is the aggregate production technology, the second is the resource constraint, the third is the intratemporal Euler equation at the aggregate level, the fourth is the definition of lump-sum taxes and transfers given our assumptions about firm ownership and fiscal policy, the fifth is the level of government purchases given our assumptions about government spending, and the sixth equation is the monetary policy rule. It is worth emphasizing that we could make different assumptions about how agents form beliefs and which of these static relationships they know with certainty (if any). Define $a_t = [C_t, \pi_t, Y_t, N_t, G_t, w_t, \tau_t, R_t]'$ and note that the six equations 49-54 uniquely determine a_t given x_t.[30] These static relationships are (log-)linearized to be

$$\widehat{Y}_t = \widehat{N}_t \tag{55}$$

$$Y\widehat{Y}_t = \left(1 + \frac{\Phi}{2}(\pi - 1)^2\right)\left(C\widehat{C}_t + G\widehat{G}_t\right) + (C + G)\,\Phi\,(\pi - 1)\,\pi\widehat{\pi}_t \tag{56}$$

$$\widehat{w}_t = \widehat{N}_t + \widehat{C}_t \tag{57}$$

$$\begin{aligned}\tau\widehat{\tau}_t =& G\widehat{G}_t - (1 - w)Y\widehat{Y}_t + wY\widehat{w}_t + \Phi(\pi - 1)\pi(C + G)\widehat{\pi}_t \\ &+ \frac{\Phi}{2}(\pi - 1)^2\left(C\widehat{C}_t + G\widehat{G}_t\right)\end{aligned} \tag{58}$$

$$\widehat{G}_t = 0 \tag{59}$$

$$\widehat{R}_t = \frac{\alpha_\pi}{R}\widehat{\pi}_t + \frac{\alpha_Y}{R}\widehat{Y}_t. \tag{60}$$

If we are linearizing around a steady state in which $R_t = 1$, then equation 78 is replaced by $\widehat{R}_t = 0$, which holds locally to that steady state. The six equations 55-78 take the form of

[30] To see this, note that $G_t = G$ is known to households and firms. With knowledge of C_t and π_t, they can compute Y_t from equation 50. Then the value of N_t follows from equation 49. Then the value of w_t follows from equation 51. Finally, the values of τ_t and R_t follow from equations 52 and 54.

those in the vector equation 38.

The household flow budget constraint and optimality conditions are (log-)linearized to be

$$(C + wN)\,\widehat{C}_{h,t} = \frac{1}{\pi}\widehat{b}_{h,t-1} - \frac{1}{R}\widehat{b}_{h,t} + 2wN\widehat{w}_t - \tau\widehat{\tau}_t \tag{61}$$

$$\widehat{C}_{h,t} - R\psi\widehat{b}_{h,t} = -\widehat{R}_t + \mathbb{E}_t\left[\widehat{C}_{h,t+1} + \widehat{\pi}_{t+1}\right]. \tag{62}$$

The firm's optimality condition is log-linearized to be

$$\begin{aligned}
&Y\left(1+\varepsilon(w-1)\right)\widehat{p}_{f,t} + (1-w)Y\widehat{Y}_t - Yw\widehat{w}_t \\
&+\frac{\Phi}{\varepsilon-1}\pi\left\{(C+G)\,(2\pi-1)\,\widehat{\pi}_{f,t} + (\pi-1)\left(C\widehat{C}_t + G\widehat{G}_t\right)\right\} = \\
&\beta\frac{\Phi}{\varepsilon-1}\pi(\pi-1)(C+G)\left(\widehat{C}_t - \mathbb{E}\widehat{C}_{t+1}\right) \\
&+\beta\frac{\Phi}{\varepsilon-1}\pi\left\{(C+G)\,(2\pi-1)\,\mathbb{E}_t\widehat{\pi}_{f,t+1} + (\pi-1)\left(C\mathbb{E}_t\widehat{C}_t + G\mathbb{E}_t\widehat{G}_{t+1}\right)\right\}
\end{aligned} \tag{63}$$

$$\widehat{\pi}_{f,t} = \widehat{p}_{f,t} - \widehat{p}_{f,t-1} + \widehat{\pi}_t \tag{64}$$

Define $h_t = [C_{h,t}, b_{h,t}, \pi_{f,t}, p_{f,t}]'$ and $e_t = [\mathbb{E}_t a'_{t+1}, \mathbb{E}_t h'_{t+1}]'$. Note that equations 61, 62, 63, and 64 take the form of those in vector equation 39. Define $y_t = [h'_t, a'_t, e'_t]'$. Using the vector equation 40 (with the definitions of a_t and e_t used here), we can put the system in the form of equation 41.

Note that y_t has 24 elements. There are two household-specific equations, two firm-specific equation, two equations regarding beliefs, six static relationships mapping x_t to the rest of a_t, and twelve equations stemming from the vector equation 40. So, we have 24 equations when the system in written in the form of equation 41.

With a solution to the (log-)linear approximation to the equilibrium, we require that the perceived values of $\nu_{C,t}$ and $\nu_{\pi,t}$ induce the following outcomes

$$C_{h,t} = C_t \tag{65}$$

$$\pi_{f,t} = \pi_t. \tag{66}$$

Equation 65 ensures that the bond market clears because, along with the budget constraint, it ensures that $b_{h,t} = 0$. Equation 66 ensures that the goods market clears because it ensures that the competitive firm that aggregates goods from the monopolists earns zero profits.[31]

[31] Note that we imposed labor maket clearing by imposing the intratemporal Euler equation at both the household and aggregate levels.

These restrictions take the form of those in equation 45. We use these restrictions to solve for the perceived values of $\nu_{C,t}$ and $\nu_{\pi,t}$ that will lead to market clearing, which delivers the period equilibrium.

D.2 The period equilibrium in the simple NK model with only sticky prices at the ZLB

Here, we describe how we construct the period equilibrium in the simple NK model with only sticky prices when the ZLB is binding. Here, we maintain the assumption that households and firms know the REE in the steady state. As a result, households and firms know the decision rules they will use in the event that the economy reverts to steady state, and those rules take the form

$$\widehat{b}_{h,t} = \omega_{b,b}\widehat{b}_{h,t-1} \tag{67}$$

$$\widehat{C}_{h,t} = \omega_{C,b}\widehat{b}_{h,t-1} \tag{68}$$

$$\widehat{p}_{f,t} = \omega_{p,p}\widehat{p}_{f,t-1} \tag{69}$$

$$\widehat{\pi}_{f,t} = \omega_{\pi,p}\widehat{p}_{f,t-1}. \tag{70}$$

To fix notation, we let $\tilde{X}$ be the REE value of the variable X_t when the ZLB binds and let $\widehat{\tilde{X}}_t = \log(X_t/\tilde{X})$. Note that we continute to use $\widehat{b}_{h,t}$ and $\widehat{p}_{f,t}$ at the ZLB, because the REE values of those variables are equal to the same values at the ZLB and in steady state.

Households and firms form beliefs using

$$\widehat{\tilde{C}}_t = \tilde{m}_{C,t-1} + \tilde{\sigma}_{\nu,C,t-1}\tilde{\nu}_{C,t} \tag{71}$$

$$\widehat{\tilde{\pi}}_t = \tilde{m}_{\pi,t-1} + \tilde{\sigma}_{\nu,\pi,t-1}\tilde{\nu}_{\pi,t} \tag{72}$$

Households and firms know that equations 49-54 hold, and we log-linearize those equations around the REE values while the ZLB binds to get

$$
\begin{aligned}
\widehat{\tilde{Y}}_t =& \widehat{\tilde{N}}_t && (73)\\
\tilde{Y}\widehat{\tilde{Y}}_t =& \left(1+\frac{\Phi}{2}\left(\tilde{\pi}-1\right)^2\right)\left(\tilde{C}\widehat{\tilde{C}}_t+\tilde{G}\widehat{\tilde{G}}_t\right)+\left(\tilde{C}+\tilde{G}\right)\Phi\left(\tilde{\pi}-1\right)\tilde{\pi}\widehat{\tilde{\pi}}_t && (74)\\
\widehat{\tilde{w}}_t =& \widehat{\tilde{N}}_t+\widehat{\tilde{C}}_t && (75)\\
\tilde{\tau}\widehat{\tilde{\tau}}_t =& \tilde{G}\widehat{\tilde{G}}_t-(1-\tilde{w})\tilde{Y}\widehat{\tilde{Y}}_t+\tilde{w}\tilde{Y}\widehat{\tilde{w}}_t\\
&+\Phi(\tilde{\pi}-1)\tilde{\pi}(\tilde{C}+\tilde{G})\widehat{\tilde{\pi}}_t+\frac{\Phi}{2}(\tilde{\pi}-1)^2\left(\tilde{C}\widehat{\tilde{C}}_t+\tilde{G}\widehat{\tilde{G}}_t\right) && (76)\\
\widehat{\tilde{G}}_t =& 0 && (77)\\
\widehat{\tilde{R}}_t =& 0. && (78)
\end{aligned}
$$

The household flow budget constraint and optimality conditions are (log-)linearized to be

$$
\begin{aligned}
\left(\tilde{C}+\tilde{w}\tilde{N}\right)\widehat{\tilde{C}}_{h,t} =& \frac{1}{\tilde{\pi}}\widehat{b}_{h,t-1}-\frac{1}{\tilde{R}}\widehat{b}_{h,t}+2\tilde{w}\tilde{N}\widehat{\tilde{w}}_t-\tilde{\tau}\widehat{\tilde{\tau}}_t && (79)\\
\widehat{\tilde{C}}_{h,t}-\tilde{R}\psi\widehat{b}_{h,t} =& \tilde{R}\frac{\tilde{\beta}}{\tilde{\pi}}p\mathbb{E}_t\left[\widehat{\tilde{C}}_{h,t+1}+\widehat{\tilde{\pi}}_{t+1}\right]+(1-p)\frac{\tilde{R}\tilde{\beta}\tilde{C}}{C\pi}\omega_{C,b}\widehat{b}_{h,t}. && (80)
\end{aligned}
$$

The firm's optimality condition is (log-)linearized to be

$$
\begin{aligned}
&\tilde{Y}\left(1+\varepsilon(\tilde{w}-1)\right)\widehat{p}_{f,t}+(1-\tilde{w})\tilde{Y}\widehat{\tilde{Y}}_t-\tilde{Y}\tilde{w}\widehat{\tilde{w}}_t\\
&+\frac{\Phi}{\varepsilon-1}\tilde{\pi}\left\{\left(\tilde{C}+\tilde{G}\right)\left(2\tilde{\pi}-1\right)\widehat{\tilde{\pi}}_{f,t}+\left(\tilde{\pi}-1\right)\left(\tilde{C}\widehat{\tilde{C}}_t+\tilde{G}\widehat{\tilde{G}}_t\right)\right\}=\\
&p\tilde{\beta}\frac{\Phi}{\varepsilon-1}\tilde{\pi}(\tilde{\pi}-1)(\tilde{C}+\tilde{G})\left(\widehat{\tilde{C}}_t-\mathbb{E}\widehat{\tilde{C}}_{t+1}\right)\\
&+p\tilde{\beta}\frac{\Phi}{\varepsilon-1}\tilde{\pi}\left\{\left(\tilde{C}+\tilde{G}\right)\left(2\tilde{\pi}-1\right)\mathbb{E}_t\widehat{\tilde{\pi}}_{f,t+1}+\left(\tilde{\pi}-1\right)\left(\tilde{C}\mathbb{E}_t\widehat{\tilde{C}}_t+\tilde{G}\mathbb{E}_t\widehat{\tilde{G}}_{t+1}\right)\right\}\\
&+(1-p)\tilde{\beta}\frac{\tilde{C}}{C}\frac{\Phi}{\varepsilon-1}\pi\left\{\left(C+G\right)\left(2\pi-1\right)\mathbb{E}_t\omega_{\pi,p}\widehat{p}_{f,t}\right\}\\
&+(1-p)\tilde{\beta}\frac{\Phi}{\varepsilon-1}\pi(\pi-1)(C+G)\frac{\tilde{C}}{C}\widehat{\tilde{C}}_t && (81)\\
&\widehat{\tilde{\pi}}_{f,t}=\widehat{p}_{f,t}-\widehat{p}_{f,t-1}+\widehat{\tilde{\pi}}_t && (82)
\end{aligned}
$$

We can now proceed as in the previous subsection to solve the linear system.

We need an approximation in which households and firms contemplate variables at the ZLB that grow slower than $\tilde{R}^{-1}\tilde{\pi}\tilde{\beta}^{-1}p^{-1}$. This growth condition keeps $\widehat{\tilde{C}}_{h,t}$ bounded in equation 80 and keeps $\widehat{\tilde{\pi}}_{f,t}$ bounded in equation 81. Because of the probability of switching from the ZLB to the steady state, the growth condition also ensures that the unconditional expectation of each variable in the system converges to zero as the horizon of the expectation goes to infinity. The growth condition is satisfied by only one solution to the linear system, and we study that solution. That solution to the linear system is similar to the non-linear solution to the model. Note that in any learning equilibrium, $\widehat{b}_{h,t} = \widehat{p}_{f,t} = 0$ in every period. So the learning equilibrium is always non-explosive.

E A comparison to Euler equation learning

An alternative approach to learning is to start with aggregate Euler equations from the REE. This approach has been used by Evans and Honkapohja (2006) and Ferrero (2007). The Euler equation approach to learning stands in contrast to the approach of Preston (2005) and Eusepi et al. (2022) in which households and firms solve their infinite horizon problems given beliefs. Our approach follows Preston (2005) and Eusepi et al. (2022). In this appendix, we analyze b under the Euler equation approach to learning to determine if the way learning is modeled has meaningful implications for the speed of convergence.

In our simple NK model with only sticky prices, the aggregate Euler equations from the REE are given by

$$\widehat{C}_t = -\beta\alpha_\pi\widehat{\pi}_t + \mathbb{E}_t\left[\widehat{C}_{t+1} + \widehat{\pi}_{t+1}\right]$$
$$\widehat{\pi}_t = \frac{2(\varepsilon-1)}{\Phi}\widehat{C}_t + \beta\mathbb{E}_t\widehat{\pi}_{t+1}.$$

Under the assumption that the learning is about steady states (or means), the Euler equation learning approach proceeds by assuming that $\mathbb{E}_t\widehat{C}_{t+1} = m_{C,t-1}$ and $\mathbb{E}_t\widehat{\pi}_{t+1} = m_{\pi,t-1}$ so that

$$\widehat{C}_t = -\beta\alpha_\pi\widehat{\pi}_t + m_{C,t-1} + m_{\pi,t-1}$$
$$\widehat{\pi}_t = \frac{2(\varepsilon-1)}{\Phi}\widehat{C}_t + \beta m_{\pi,t-1}.$$

The values of $\widehat{\pi}_t$ and $\widehat{C}_t$ are calculated from these equations. Beliefs are updated so that

$$m_{x,t} = m_{x,t-1} + \frac{1}{\lambda_0 + t}\left(\hat{x}_t - m_{x,t-1}\right).$$

This updating equation (in vector form) has the same form as equation (1) in the main text of our paper. When we compute the value of b under the Euler equation learning approach, we get $b = 0.95$ and $T_{2/3} \approx 1.6$ billion. Under the learning approach we take in the text, and as reported in Table 1, when we assume households and firms solve their infinite horizon problems given beliefs we get $b = 0.25$ and $T_{2/3} = 5$. We conclude that modeling agents infinite-horizon decision problems, rather than using aggregate Euler equations from the REE, can have important implications for conclusions regarding the speed of learning.

F Calvo-style nominal rigidities

In the main text of our paper we study a model with Rotemberg-style nominal rigidities. It is well known that Calvo-style nominal rigidities are equivalent to Rotemberg-style nominal rigidities in an REE when linearizing a model around the target-inflation steady state. However, these models need not be equivalent with analyzing learning equilibria. An advantage of using Rotemberg-style nominal rigidities is that in the fully nonlinear model it does not have the price-dispersion state variables that appear in models with Calvo-style nominal rigidities. Because it is more-feasible to solve the nonlinear model with fewer state variables, we used Rotemberg-style nominal rigidities. In this appendix, we re-do our analysis for the simple NK model using Calvo-style nominal rigidities and show that our conclusions are robust. We present details of the model in the sub-section below. Table 4 reports the values of b in the Rotemberg and Calvo versions of the model.

Table 4: b in the simple NK model

	b in the Rotemberg model	b in the Calvo model
Sticky prices and wages	0.63	0.82
Only sticky prices	0.25	0.26
Only sticky wages	0.58	0.82
Higher α_π ($\alpha_\pi = 3$)	0.18	0.47
Lower α_π ($\alpha_\pi = 1.01$)	0.98	0.97
Lower α_Y ($\alpha_Y = 0$)	-0.30	0.00

Source: Authors' calculations.