

NBER WORKING PAPER SERIES

SLOW LEARNING

Lawrence Christiano
Martin S. Eichenbaum
Benjamin K. Johannsen

Working Paper 32358
<http://www.nber.org/papers/w32358>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2024, Revised June 2026

The analysis and conclusions set forth are those of the authors and do not indicate concurrence by the Board of Governors, Federal Reserve staff, the National Bureau of Economic Research, 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, Harold Uhlig, Ivan Werning, Fabian Winkler, Mike Woodford, and participants in conferences and seminars for helpful comments and suggestions.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w32358>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Lawrence Christiano, Martin S. Eichenbaum, and Benjamin K. Johannsen. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Slow Learning

Lawrence Christiano, Martin S. Eichenbaum, and Benjamin K. Johannsen

NBER Working Paper No. 32358

April 2024, Revised June 2026

JEL No. E12, E39, E70

ABSTRACT

This paper provides an analytic characterization of the speed of convergence under learning to a rational expectations equilibrium (REE) for a large class of multivariate models. We show that learning is slower when people's beliefs about model outcomes are more self-fulfilling. The paper also investigates which features of a model economy make beliefs more self-fulfilling, using variants of the simple new-Keynesian model and a medium-scale DSGE model. For empirically plausible specifications of these models, convergence of a learning equilibrium to the REE is so slow that analysis based on rational expectations can be misleading.

Lawrence Christiano
Northwestern University
Department of Economics
and NBER
l-christiano@northwestern.edu

Benjamin K. Johannsen
Federal Reserve Board of Governors
benjamin.k.johannsen@frb.gov

Martin S. Eichenbaum
Northwestern University
Department of Economics
and NBER
eich@northwestern.edu

1 Introduction

The founders of the rational expectations (RE) revolution thought that, to be useful, rational expectations equilibria (REE) must be the outcome of learning equilibria.¹ While convergence to an REE may be a necessary condition for rational expectations to be a useful model of expectations, it is not sufficient. As Vives (1993) stresses, in a changing world, for all practical purposes, “‘slow’ convergence may mean no convergence.” The critical issue is whether the speed of convergence is fast enough to render RE a useful guide for normative and positive analyses. In this paper, we characterize, for a large class of models, the speed of convergence of a learning equilibrium to an REE.

We depart from the RE framework and assume that people learn about their environment by forming beliefs about future economic outcomes and updating those beliefs as data arrive. We analytically characterize the speed of convergence in a broad class of non-stochastic learning models. 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 demonstrate that our results can be applied to a broad class of macroeconomic models that are solved using linearization methods. This class includes medium-scale dynamic stochastic general equilibrium (DSGE) models with shocks, like Christiano et al. (2005) (hereafter, CEE).

An important contribution of this paper is to provide a formal characterization of the asymptotic rate of convergence for a broad class of nonlinear, non-stochastic multivariate learning models studied in Evans and Honkapohja (2000). Proposition 1 of this paper shows that a particular scalar statistic, b , of the system determines the asymptotic speed of convergence to an REE. That parameter, which also determines the E-stability of an REE ($b < 1$), can be calculated from the model solution (see Evans and Honkapohja (2000) for a discussion of E-stability). A model exhibits slow learning when b is less than but close to unity. We investigate the economic determinants of b , i.e., whether learning is fast or slow.

Our central finding is that, in empirically plausible new-Keynesian (NK) models, learning equilibria converge slowly to REE. Indeed, learning can be extraordinarily slow. We develop an asymptotically valid formula that depends only on b which approximates the time required to close a given fraction of the gap between initial and RE beliefs. Using simulations, we show that our formula is quite accurate for all the models considered. In the simple NK model,

¹For example, in his seminal paper on asset pricing, Lucas (1978) writes ‘....the model described above “assumes” that people know a great deal about the structure of the economy, and perform some non-routine computations. It is in order to ask, then: will an economy with people armed with “sensible” rules-of-thumb, revising these rules from time to time so as to claim observed rents, tend as time passes to behave as described....’

when wages and prices are sticky, progress to the REE is measured in decades.² An even stronger result holds for the benchmark DSGE model developed in Christiano et al. (2005): meaningful convergence to the REE is measured in *centuries*. Even more dramatically, in the simple NK model, when the zero lower bound (ZLB) is binding, progress is measured in *millennia*. We argue that stickiness in wages and the parameters of the monetary policy are the critical determinants of the speed of convergence.

The intuition behind our results is that people’s beliefs influence actual economic outcomes. For example, suppose people revise their expectations of future inflation upward. Wage setters, recognizing they would like higher nominal wages in the future, have an incentive to front-load those increases when wage stickiness is present. By increasing nominal marginal cost, this rise in the current wage exerts upward pressure on current inflation. In effect, sticky wages make inflation expectations at least partially self-fulfilling, much like in the wage-price spiral literature (see Blanchard (1986)). The more self-fulfilling expectations are, the slower people are to abandon initial priors (the higher is b) and the longer it takes for expectations to converge. Monetary policy also plays an important role. The more (less) aggressive the monetary authority’s response to inflation (the output gap), the less likely it is that inflation expectations will be self-fulfilling.

Sections 5 and 6 analyze the robustness of our key conclusions to a variety of perturbations. First, as noted above, those conclusions hold in empirically plausible DSGE models like the model of Christiano et al. (2005).³ Second, we consider the speed of learning when nominal rigidities arise from Rotemberg rather than Calvo-style nominal rigidities. Under RE, Calvo-style and Rotemberg-style nominal rigidities are equivalent in a first-order approximation to the model.⁴ However, they need not be equivalent in a learning equilibrium. As it turns out, our key results about the speed of learning are robust across the two different ways of modeling nominal rigidities. Third, our results are robust to whether or not the zero lower bound (ZLB) on the nominal interest rate is binding. Fourth, we show that our results are robust to whether people solve their problems using a version of Kreps (1998)’s *anticipated utility* (our benchmark analysis) or *internally rational* learning. In the anticipated utility framework, people update their beliefs at each period as new data comes in. But when they make their decisions, they proceed as though their beliefs will never be revised.⁵ Under internally rational learning, people fully integrate the fact that they are learning when they

²In this model there is no capital, no habit persistence in consumption, and a static Taylor rule.

³The nonlinear model underlying Smets and Wouters (2007) is the same as Christiano et al. (2005) up to the specification of shocks, so our DSGE results also apply to that model. This is consistent with our conclusion that learning is slow in estimated, empirically plausible NK models.

⁴See, for example, Born and Pfeifer (2020).

⁵This approach has been criticized for its internal inconsistency (see Adam and Marcet (2011)). See also, Adam et al. (2017), Adam and Merkel (2019), Winkler (2020), and Caines and Winkler (2021).

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 model’s nonlinear solution is computationally challenging.⁶ The associated computational burden explains why much of the learning literature works with a version of Kreps (1998)’s *anticipated utility* approach. Because of computational challenges, we investigate the implications of the two different learning models using the full nonlinear solution to the simple NK model when the ZLB is binding. Fifth, we demonstrate that our results are robust to whether we work with the fully nonlinear 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 Christiano et al. (2005). So we use the simple NK model when the ZLB is binding. In the linear and nonlinear versions of that model, the speed of convergence to the REE is very slow.

Section 7 analyzes the impact of learning on how fiscal and monetary policies impact the economy. We show that the speed of convergence plays a crucial role in assessing the efficacy of government policies, such as increases in government purchases. These results are most stark in the simple NK model when the ZLB is binding.⁷ We make this case using a nonlinear version of the NK model in which only prices are sticky. This model has been widely used in the literature that analyzes the effect government purchases when the ZLB is binding. In the model we consider, the effects of a rise in government purchases are much smaller under learning than under RE. Under RE, 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. In the learning model, expected inflation is partially backward-looking and changes little after an increase in government purchases. So, the real interest rate does not 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. Consistent with the existing literature (for example, Farhi and Werning (2019), Del Negro et al. (2023), and Woodford (2012)), we find that even this simple form of forward guidance is powerful under RE. As is well known, the power of forward guidance under RE reflects its strong effect on expected inflation. Under learning, expectations are partially

⁶See Appendix F for details.

⁷See, for example, Eggertsson and Woodford (2003), Christiano et al. (2011), and Del Negro et al. (2023).

backward-looking, and forward guidance is not very powerful. So, as with fiscal policy, an REE-based analysis of monetary policy can be very misleading.

The models we work with have multiple REEs. Benhabib et al. (2001b) document that NK models have two REE steady states. In one steady state the ZLB is binding, in the other it is not. Mertens and Ravn (2014) and Bilbiie (2022) discuss ZLB REEs that are caused by nonfundamental shocks. Throughout our paper, when we use linear approximations of the model, we focus on learning equilibria that are local to E-stable REEs.⁸ We do so for three reasons. First, E-stable REEs are the most-commonly studied in the related literature. Second, a full treatment of equilibrium multiplicity in the NK model is beyond the scope of this paper. Third, we find that learning can be very slow, even in learning equilibria that are local to E-stable REEs. This finding poses a challenge to policy analyses based on E-stable REEs that are the focus of the literature. We discuss equilibrium multiplicity in Appendix G.

The remainder of this paper is organized as follows. Section 2 discusses the relationship of our paper to the related literature. Section 3 presents our theoretical results regarding the speed of convergence to an REE. Section 4 analyzes the speed of convergence in the simple NK model when the ZLB is not binding. Section 5 analyzes the speed of convergence in the CEE model when the ZLB is not binding. Section 6 discusses the robustness of our analysis, including linear and nonlinear versions of the simple NK model when the ZLB is binding, and versions of the model under anticipated utility and internally rational learning. Section 7 considers whether learning matters for policy analysis. Finally, section 8 contains concluding remarks.

2 Related Literature

Our paper is related to several literatures. First, a large literature studies the properties of recursive estimators in learning models with decreasing gain.⁹ Ljung (1977) establishes almost sure convergence. 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. In contrast, we study the speed at which learning occurs. Marcet and Sargent (1995) study the speed of convergence using numerical

⁸We study minimum-state-variable equilibria. McCallum (1983) offers arguments in favor of focusing on minimum-state-variable equilibria. See Arifovic et al. (2018) for a different approach to learning under which the ZLB steady state is stable under learning. See Eusepi (2005) for a discussion of learning equilibria that converge to cycles.

⁹Another branch of the literature studies the properties of recursive estimators in learning models with constant gain. See, for example, Marcet and Nicolini (2003). We focus on decreasing gain because it arises naturally in the context of Bayesian learning about parameters.

methods. In contrast, we provide analytic results for a large class of models. In addition, our numerical results pertain to a class of empirically plausible DSGE models.

We work in a nonstochastic environment in which it is natural to focus on the rate of convergence of deterministic sequences. The stochastic approximation literature focuses on convergence in distribution. Benveniste et al. (1990) show that $t^{1/2}\theta_t$ converges to a Normal distribution with mean zero for $b < 1/2$.¹⁰ Christopheit and Massmann (2018) provide a derivation of the speed of convergence in distribution of a learning equilibrium for $b \in [1/2, 1)$. They do so for a scalar linear regression model. Their theoretical results do not, in general, hold for nonlinear, multivariate environments. We provide multivariate, nonstochastic results on the speed of convergence of θ_t for $b < 1$ and apply them to the nonlinear NK model.

A third literature examines the speed of convergence in the simple NK model (see, e.g., Ferrero (2007)). We study rates of convergence in a range of models, including an empirically plausible medium-sized DSGE model. Additionally, Ferrero (2007) uses the Euler equation approach to learning.¹¹ We consider that approach to learning as well as the anticipated utility and internally rational approaches. In contrast to Ferrero (2007), we also study convergence rates in linear and nonlinear versions of the NK model when the ZLB is binding.

Heemeijer et al. (2009) and Hommes (2011) study positive and negative feedback loops from expectations to outcomes using laboratory experiments and univariate models with constant gain. Their results provide experimental support for the class of learning models considered in this paper.

A different literature assumes that people have RE and studies how people form beliefs about exogenous variables when they have imperfect information. For example, Vives (1993) asks: how quickly do people’s beliefs about an exogenous cost parameter converge? Other authors like Erceg and Levin (2003), Gust et al. (2018), and Farmer et al. (2024) study how people draw inferences about the movements in the unobserved transitory and persistent components of an exogenous variable. In contrast, we study convergence rates for beliefs about objects whose values depend on those beliefs. Other papers that study the dynamic

¹⁰Evans and Honkapohja (2001) discuss the speed of convergence in Chapter 7 of their book. They cite Benveniste et al. (1990) regarding the rate of convergence of the stochastic learning algorithm to a non-degenerate random variable for the case of $b < 1/2$. Evans and Honkapohja (2001) return to the speed of convergence in Chapter 15.2. There, they cite Marcet and Sargent (1995) for simulation results that imply learning can be very slow when $[1/2 < b < 1)$.

¹¹See Evans (2021) for a discussion of the Euler-equation approach to learning and Adam and Marcet (2011) for a critique of that approach. An earlier version of our paper (Christiano et al. (2018)) followed the Euler-equation approach to learning models (not the anticipated utility or the internally rational approaches to learning). See Eusepi and Preston (2018) for a discussion of different approaches to learning. Our current approach is more natural in models where people solve infinite-horizon optimization problems and learn by Bayesian updating. Different approaches to learning can imply different speeds of learning (see Appendix C for a discussion of the speed of learning under different assumptions).

properties of Bayesian learning with incomplete or noisy information include Collin-Dufresne et al. (2016) and Angeletos and Laó (2020). In this paper, we consider learning equilibria along the lines of those studied in, for example, Eusepi and Preston (2018), and characterize the rate of convergence of beliefs.

Another literature studies the policy implications of learning when the ZLB is binding. Evans et al. (2008) study the effects of fiscal policy and alternative money supply rules for Euler-equation learning equilibria when the ZLB is binding. Eusepi (2010) shows that central bank communication can have an important impact on Euler-equation learning equilibria when the ZLB is binding. Benhabib et al. (2014) and Evans et al. (2022) study convergence in the nonlinear simple NK model with anticipated-utility learning. Eusepi et al. (2021) discuss forward guidance when the ZLB is binding in the simple NK model with anticipated utility. 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 to our analysis, these papers do not characterize rates of convergence or consider internally rational learning.

3 Learning in a Nonlinear Environment

In this section, we study the speed of convergence of learning in a non-stochastic environment. Let θ_t be a k -dimensional vector of variables that summarizes people's period- t beliefs about the values of the parameters governing a perceived data-generating process. We interpret θ_t as a deviation from a particular fixed point of beliefs in our learning algorithm.¹² As in Evans and Honkapohja (2000), the vector θ_t evolves according to

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

for $t = 1, 2, 3, \dots$ where $M(\theta_{t-1}, \gamma_t)$ is a potentially non-linear function. Here, θ_0 is given and γ_t the learning gain. Let the $k \times k$ matrix $D_1 M$ denote the derivative of M with respect to the vector θ_{t-1} and let b be the largest real part of the eigenvalues of $D_1 M(0, 0)$. Similar to Evans and Honkapohja (2000), we adopt the following two assumptions:

Assumption 1. The vector-valued function $M : \mathbb{R}^k \times \mathbb{R} \rightarrow \mathbb{R}^k$ has the following properties:

- (i) $M(0, 0) = 0$;
- (ii) M is continuously differentiable in a neighborhood of the origin;
- (iii) $M(0, \gamma_t) = 0$, for all $t \geq 1$;
- (iv) $M(0, \gamma_t)$ is continuously differentiable in a neighborhood of $(0, \gamma_t)$, for all $t \geq 1$;

¹²We do not require that this fixed point be unique.

- (v) the real parts of the eigenvalues of $D_1 M(0, 0)$ are strictly less than unity and no eigenvalue of $D_1 M(0, \gamma_t)$ is equal to $1 - \gamma_t^{-1}$ for all $t \geq 1$.

Assumption 2. The scalar series, γ_t , satisfies the following properties:

- (i) $0 < \gamma_t < 1$, for all $t \geq 1$;
- (ii) $\lim_{t \rightarrow \infty} t\gamma_t = 1$.

Our environment is the same as Evans and Honkapohja (2000)'s, except that we have a more-restrictive specification of γ_t .¹³ They establish that for $b < 1$ and θ_0 sufficiently close to 0, $\|\theta_t\|$ converges to zero. We study the rate at which $\|\theta_t\|$ converges to zero.

We adopt the following definition of the asymptotic rate of convergence of a variable in a nonstochastic environment. Our definition enables us to establish the precise sense in which the parameter b determines the rate of convergence.

Definition 1. For $b < 1$, x_t asymptotically converges to zero at the rate t^{b-1} (which we denote by $x_t \simeq t^{b-1}$) if for any $0 < \delta < 1 - b$, (i) $\lim_{t \rightarrow \infty} \frac{\|x_t\|}{t^{b-1+\delta}} = 0$, and (ii) $\lim_{t \rightarrow \infty} \frac{\|x_t\|}{t^{b-1-\delta}} = \infty$.

Here, $\|\cdot\|$ denotes a norm on $\mathbb{R}^k$. The first part of Definition 1 says that for any small, 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.

Christopeit and Massmann (2018) report results which imply that when M is linear and θ_t is a scalar then $t^{1-b} \|\theta_t\| \rightarrow \kappa$ for $0 < \kappa < \infty$. Our concept of convergence in Definition 1 is weaker than convergence to a finite, nonzero constant limit. We work with our definition because $t^{1-b} \|\theta_t\|$ may not converge to a finite, nonzero constant in nonlinear or multivariate cases.¹⁴

Conditions (i) and (ii) in Definition 1 describe a space of sequences. It is easy to verify that it includes $\theta_t = \kappa t^{b-1}$, $0 < \kappa < \infty$. It also includes sequences that cycle around t^{b-1} with decreasing amplitude as t increases, e.g. $\theta_t = [2 + \sin(t)] t^{b-1}$. In addition, it includes sequences that deviate from t^{b-1} at a slow enough rate, such as $\theta_t = \log(t) t^{b-1}$. Definition 1, in effect, defines a set of sequences that exhibit a power rate of convergence, up to slowly varying terms. Consistent with this definition, all of the elements in the set of sequences

¹³Evans and Honkapohja (2000) assume that part (i) of Assumption 2 holds. Additionally, they assume $\lim_{t \rightarrow \infty} \gamma_t = 0$ and $\lim_{t \rightarrow \infty} \sum_{j=1}^t \gamma_j = \infty$. Both of these properties are implied by part (ii) of Assumption 2.

¹⁴The following two examples illustrate this point. First, consider the linear multivariate case, $k = 2$. Let $\theta_t = [x_t, y_t]'$ and suppose $M(\theta_{t-1}, \gamma_t) = \begin{bmatrix} b & 1 \\ 0 & b \end{bmatrix} \theta_{t-1}$. By iterated substitution it can be shown that $x_t/t^{b-1} \rightarrow \infty$ as $t \rightarrow \infty$. For our second example consider the nonlinear, univariate function, $M(\theta_{t-1}, \gamma_t) = b\theta_{t-1} - \frac{\theta_{t-1}}{\log(|\theta_{t-1}|)}$ for $\theta_{t-1} \neq 0$ and $M(0, \gamma_t) = 0$. This function satisfies our assumptions, and $D_1 M(0, 0) = b$. Yet, it can be shown that $x_t/t^{b-1} \rightarrow \infty$ as $t \rightarrow \infty$. In both examples, $x_t \simeq t^{b-1}$ according to Definition 1.

identified by Definition 1 converge to zero at a rate slower than a geometric rate (i.e., the sequence λ^t , $t = 1, 2, \dots$, for $0 < \lambda < 1$).

Let U denote a neighborhood of the origin that has the property that every $\theta_0 \in U$ generates a sequence $\|\theta_t\|$ that converges to zero (Proposition 1 in Evans and Honkapohja (2000) establishes that such a U exists). We now state our proposition, which is proved in Appendix A.

Proposition 1. *Suppose θ_t evolves according to equation (1), Assumption 1 holds (which implies $b < 1$), and 2 holds. There exist uncountably many $\theta_0 \in U$ that generate a sequence $\theta_t \simeq t^{b-1}$, and for all $\theta_0 \in U$ and any $\delta > 0$, $\lim_{t \rightarrow \infty} \frac{\|\theta_t\|}{t^{b-1+\delta}} = 0$.*

For intuition, it is useful to consider the following special case: If $M(\theta_{t-1}, \gamma_t) = M\theta_t$, where M is a matrix and all the eigenvalues of M are real and distinct, then $t^{1-b}\|\theta_t\|$ converges to a non-zero limit for θ_0 in a dense subset of U .

Proposition 1 leaves open the possibility that there are values of $\theta_0 \in U$ that do not have the property, $\theta_t \simeq t^{b-1}$. An obvious, but uninteresting, example is $\theta_0 = 0$ which implies $\theta_t = 0$ for all t . A more informative example is given by:

$$\theta_t = \theta_{t-1} + \gamma_t (M - I) \theta_{t-1}, \quad M = \begin{bmatrix} b & 0 \\ 0 & a \end{bmatrix},$$

where $a < b < 1$. Suppose that $\theta_0 = [0, 1]'$, so that $\theta_t = \left[0, \prod_{j=1}^t (1 + \gamma_j (a - 1)) \right]'$. It is straightforward to show that $\theta_t \simeq t^{a-1}$, so that the speed of convergence is faster than t^{b-1} .¹⁵ However, any nonzero perturbation to the first element of θ_0 , no matter how small, gives rise to a θ_t sequence that has the property, $\theta_t \simeq t^{b-1}$. This example, in which there are special initial beliefs that converge faster than t^{b-1} , generalizes to any system in which M is linear, satisfies Assumption 1, and has an eigenvalue with real part less than b . In any such system, there are initial beliefs that are arbitrarily close to any θ_0 that give rise to a sequence $\theta_t \simeq t^{b-1}$.

In our analysis, the function, M , in equation (1) embeds the equilibrium map from beliefs, θ_{t-1} , and the gain, γ_t , to equilibrium outcomes under learning. It can be convenient to approximate equation (1) around the limit points of θ_t and γ_t using perturbation methods.¹⁶ The first-order expansion of equation (1) implies $\theta_t = \theta_{t-1}$, so that there is no learning in a neighborhood of the limit points. This property is a harbinger of the slow-learning results in this paper. A second-order expansion of equation (1) simply replaces $M(\theta_{t-1}, \gamma_t)$

¹⁵To see this, note that $\log(\theta_t)$ is approximately the partial sum of a harmonic series, $\log(\theta_t) \approx (a-1) \sum_{j=1}^t \gamma_j \approx (a-1) [\kappa + \log(t)]$, where κ is a finite constant.

¹⁶Additional smoothness assumptions about $M(\cdot, \cdot)$, beyond what appears in Assumptions 1 and 2 are required to apply perturbation methods (see Judd (1996)).

in equation (1) with its first-order expansion, $D_1 M(0, 0)\theta_{t-1}$.¹⁷ This result provides an interpretation of a standard procedure that uses a linear approximation of M when solving learning model.

4 The Speed of Learning in the Simple NK Model

In this section, we analyze the speed of learning in a simple textbook version of the NK model that incorporates a version of Kreps (1998) anticipated utility.¹⁸ Our key finding is that the degree of wage stickiness, as well as the parameters of the monetary policy rule, play a central role in the speed of learning.

4.1 The Simple NK Model with Sticky Nominal Wages and Prices

In this subsection, we present a variant of the simple NK model in which both nominal wages and prices are sticky. We model nominal price rigidities as in Calvo (1983). In the case of wage rigidities, we work with a variant of Erceg et al. (2000).

4.1.1 Firms' Problems

A final homogeneous good, Y_t , is produced by a continuum of competitive, identical firms using the technology $Y_t = \left(\int_0^1 Y_{f,t}^{\frac{\varepsilon-1}{\varepsilon}} df \right)^{\frac{\varepsilon}{\varepsilon-1}}$, where $Y_{f,t}$ denotes differentiated inputs and $\varepsilon > 1$. The representative homogeneous-good firm chooses inputs to maximize profits, $Y_t P_t - \int_0^1 Y_{f,t} P_{f,t} df$, where P_t denotes the time- t aggregate price level and $P_{f,t}$ denotes the time t price of the f^{th} good. The representative firm's first-order condition for the f^{th} input is:

$$Y_{f,t} = \left(\frac{P_{f,t}}{P_t} \right)^{-\varepsilon} Y_t. \quad (2)$$

A monopolist produces the f^{th} intermediate good with production technology $Y_{f,t} = N_{f,t}$, where $N_{f,t}$ denotes labor hired by firm f . With probability ξ_p the monopolist sets its time- t price equal to $P_{f,t-1}$. With probability $1 - \xi_p$, the monopolist resets its price optimally to

¹⁷This can be seen by taking a second order expansion of equation (1) and noting that Assumption 1 part (iii) implies $D_2 M(0, 0) = 0$. In addition, a second-order approximation to the learning model involves a second-order approximation to γ_t . This requires expressing γ_t in recursive form. In this paper, we work with a standard representation of the gain, $\gamma_t = \frac{1}{\lambda + t}$, where $\lambda > 0$. The recursive representation is given by $\gamma_t = \gamma_{t-1} (\gamma_{t-1} + 1)^{-1}$ with initial value $0 < \gamma_1 < 1$. A second-order approximation to $\gamma_t = \gamma_{t-1} (\gamma_{t-1} + 1)^{-1}$ around $\gamma_{t-1} = 0$ is given by $\hat{\gamma}_t = \hat{\gamma}_{t-1} (1 - \hat{\gamma}_{t-1})$. It can be shown that $\hat{\gamma}_t$ satisfies Assumption 2 and that $\hat{\gamma}_t$ approaches γ_t from below.

¹⁸See Eusepi and Preston (2018) for additional discussion of anticipated utility.

$\tilde{P}_{f,t}$. It chooses $\tilde{P}_{f,t}$ to maximize the expected discounted value of profits:

$$\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \xi_P^j C_{t+j}^{-1} \left(\frac{\tilde{P}_{f,t}}{P_{t+j}} - (1 - \tau_f) s_{t+j} \right) \left(\frac{\tilde{P}_{f,t}}{P_{t+j}} \right)^{-\varepsilon} Y_{t+j}. \quad (3)$$

Here, $\mathbb{E}_t$ denotes the expectation operator under monopolists' subjective beliefs at time t , s_{t+j} denotes real marginal cost in period $t+j$, and C_{t+j}^{-1} denotes the marginal value of profits in units of the consumption good in period $t+j$. Finally, τ_f is a government tax subsidy designed to eliminate the effect of monopoly distortions in the steady state.¹⁹

4.1.2 Households' Consumption Problem

There is a continuum of identical households, $h \in (0, 1)$. Each 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 nominal payoff at the beginning of period t to nominal bonds purchased in the previous period. In period t , the household chooses consumption, $C_{h,t}$, and end-of-period bond holdings, $b_{h,t}$, to maximize:

$$\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \left\{ \log(C_{h,t+j}) - \frac{\chi}{2} \int_0^1 N_{i,h,t+j}^2 di \right\} \quad (4)$$

subject to the budget constraints, for $j \geq 0$,

$$C_{h,t+j} + \frac{b_{h,t+j}}{R_{t+j}} \leq \frac{b_{h,t+j-1}}{\pi_{t+j}} + \int_0^1 (1 + \tau_w) w_{i,t+j} N_{i,h,t+j} di - \tau_{t+j}. \quad (5)$$

Here, $\pi_t = P_t/P_{t-1}$, τ_t denotes lump-sum taxes net of profits from the firms, and R_t denotes the nominal rate of interest. The h^{th} household has a variety of members, indexed by $i \in (0, 1)$, with imperfectly substitutable types of labor, $N_{i,h,t}$. The household takes $w_{i,t}$ and $N_{i,h,t}$ as given, for $i \in (0, 1)$. We assume households have the same expectation operator, $\mathbb{E}_t$, as monopolists. For reasons similar to those discussed in Adam and Marcet (2011) we place an upper and lower bound on $b_{h,t}$. When we work with the linearized version of the model we mimic the effects of these bounds by adding the term, $\frac{\vartheta}{2} b_{h,t+j}^2$, to the left side of the household budget constraint, equation (5). In practice, we set ϑ to be a very small number and no costs are actually paid (even in learning equilibria) because $b_{h,t}$ must be zero for bond markets to clear.

4.1.3 Employment and Wage Determination

Corresponding to each $i \in (0, 1)$ there is a single labor union. For every household, h , the i^{th} member belongs to that union. The union sets the nominal wage, $W_{i,t}$, subject to Calvo

¹⁹This subsidy satisfies $(1 - \tau_f) \varepsilon / (\varepsilon - 1) = 1$.

frictions. With probability $1 - \xi_W$, union i can reset $W_{i,t}$ and we denote its choice by $\tilde{W}_t$. Otherwise the union must set $W_{i,t} = W_{i,t-1}$. Union i provides all labor demanded in equal amounts from each household. Households must supply whatever amount of the i^{th} type of labor that union i requires.

A competitive, representative firm uses the labor, $N_{i,t} = \int_0^1 N_{i,h,t} dh$, supplied by union $i \in (0, 1)$, to produce a homogeneous labor service, N_t , using the technology $N_t = \left(\int_0^1 N_{i,t}^{\frac{\epsilon-1}{\epsilon}} di \right)^{\frac{\epsilon}{\epsilon-1}}$, $\epsilon > 1$. The firm takes $W_{i,t}$ as given and sells N_t at the price, W_t . In equilibrium, $W_t = \left(\int_0^1 W_{i,t}^{1-\epsilon} di \right)^{\frac{1}{1-\epsilon}}$. Profit maximization by the competitive, representative firm implies:

$$N_{i,t} = \left(\frac{W_{i,t}}{W_t} \right)^{-\epsilon} N_t. \quad (6)$$

The wage, $\tilde{W}_{i,t}$, is chosen to maximize:

$$\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \xi_W^j \left\{ \frac{C_{t+j}^{-1}}{P_{t+j}} (1 + \tau_w) \tilde{W}_{i,t} \left(\frac{\tilde{W}_{i,t}}{W_{t+j}} \right)^{-\epsilon} N_{t+j} - \frac{\chi}{2} \left(\left(\frac{\tilde{W}_{i,t}}{W_{t+j}} \right)^{-\epsilon} N_{t+j} \right)^2 \right\}. \quad (7)$$

We assume unions have the same expectation operator, $\mathbb{E}_t$, as households and monopolists. In an REE, maximizing equation (7) corresponds to maximizing the utility of the representative household.

4.1.4 Monetary, Fiscal Policy, and Aggregation

Aggregate output is given by:

$$Y_t = C_t + G_t. \quad (8)$$

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

$$R_t = \beta^{-1} + \alpha_\pi (\pi_t - 1) + \alpha_Y \frac{Y_t - Y}{Y}. \quad (9)$$

Here, Y is the value of Y_t in the target-inflation ($\pi_t = 1$) REE steady state. The government finances subsidies to households and firms with lump-sum taxes and balances its budget each period. The government also purchases final goods, $G_t = G$.

The firm production function implies that gross output is given by:

$$Y_t = p_t^* N_t, \quad (10)$$

where p_t^* is a measure of price dispersion given by $(p_t^*)^{-1} = \int_0^1 \left(\frac{P_{f,t}}{P_t} \right)^{-\epsilon} df$ (see Yun (1996)).

4.1.5 The Evolution of Beliefs Under Anticipated Utility

For the period- t problems of households, unions, and firms to be well-defined, they must have beliefs about the aggregate variables that determine their payoffs:

$$a_{t+j} = \left[C_{t+j}, \pi_{t+j}, R_{t+j}, Y_{t+j}, N_{t+j}, w_{t+j}, \tau_{t+j}, p_{t+j}^*, G_{t+j}, s_{t+j} \right]', \quad (11)$$

where $w_{t+j} = W_{t+j}/P_{t+j}$, for $j = 0, 1, 2, \dots$. We describe a recursive mechanism that characterizes people's beliefs about $\{a_{t+j}\}_{j=0}^{\infty}$.

To focus attention on learning about the data generating process of key variables that are determined by intertemporal decisions, we work with a smaller set of variables, x_{t+j} :

$$x_{t+j} = \left[C_{t+j}, \pi_{t+j}, w_{t+j} \right]'. \quad (12)$$

We assume people know seven aggregate equilibrium conditions of the model which map from a_{t+j-1} and x_{t+j} to the non- x_{t+j} elements of a_{t+j} .²⁰ We denote this mapping, together with the relevant identities, by F :

$$a_{t+j} = F(x_{t+j}, a_{t+j-1}). \quad (13)$$

Before people observe x_t , their beliefs about x_{t+j} are summarized by a density function, D , parameterized by a finite dimensional vector Θ_{t-1} :

$$x_{t+j} \sim D(a_{t+j-1}; \Theta_{t-1}), \quad (14)$$

for $j = 0, 1, \dots$. The presence of Θ_{t-1} in equation (14) for each $j \geq 0$ reflects the anticipated utility assumption. Under this assumption, people making time- t decisions act as if their beliefs, Θ_{t-1} , will remain unchanged in the future, regardless of the data they see.

Equations (13) and (14) provide a recursive representation of people's beliefs about $\{a_{t+j}\}_{j=0}^{\infty}$ in period t before observing a_t . These beliefs are conditional on a_{t-1} and Θ_{t-1} .

After time- t decisions are made and markets clear, people use the observed x_t to update their beliefs, Θ_{t-1} , as follows:

$$\Theta_t = L(\Theta_{t-1}, x_t, a_{t-1}; \gamma_t). \quad (15)$$

Under anticipated utility, people making their decisions in period t do not internalize that they will update their beliefs at every future date according to equation (15). The form of L depends on the assumed model of learning and decision making. The variable γ_t is a

²⁰Specifically, people believe that $s_t = w_t$, $G_t = G$, $(p_t^*)^{-1} = (1 - \xi_P) \left(\frac{1 - \xi_P \pi_t^{\varepsilon-1}}{1 - \xi_P} \right)^{-\frac{\varepsilon}{1-\varepsilon}} + \xi_P \pi_t^{\varepsilon} (p_{t-1}^*)^{-1}$, $\tau_t = G_t + \tau_w w_t N_t - \left(1 - \frac{w_t}{p_t^*}\right) Y_t$ and that equations (8), (9), and (10) hold. These are equilibrium conditions in REE and, by construction, in our learning equilibria.

decreasing gain and, in the models that we consider, takes the form $\gamma_t = (\lambda + t)^{-1} > 0$.

4.1.6 Equilibrium

At the beginning of period t people observe the aggregate state, a_{t-1} . People make their decisions, conditional on their individual beliefs, Θ_{t-1} , based on observing x_t , which they interpret as a draw from D in equation (14) for $j = 0$. They compute the remaining current period aggregates, a_t , using equation (13). When people solve their period- t problems, beliefs about the distribution of future aggregates, $\{a_{t+j}\}_{j=1}^{\infty}$, are constructed using equations (13) and (14). When aggregated, all these decisions give rise to an actual set of period- t aggregates. A period t equilibrium is a fixed point of this mapping so that the observed x_t is consistent with the decisions people make. Formally,

Definition 2. Given Θ_{t-1} and a_{t-1} , a period- t equilibrium is a set of values of a_t such that:

(i) Firms, households, and unions solve their time- t optimization problems by maximizing equations (3), (4), and (7), respectively, given the mappings in equations (13) and (14); monetary policy satisfies equation (9); $G_{t+j} = G$ for $j \geq 0$, and the government balances its budget with lump-sum taxes.

(ii) Time- t bond, labor, and goods markets clear.

The seven aggregate equilibrium conditions used to construct F in equation (13) and the three market-clearing conditions in (ii) determine the values of the ten period- t equilibrium variables in a_t .²¹

Definition 3. Conditional on Θ_0 and a_0 , a *learning equilibrium* is a sequence, $t = 1, \dots$, of period equilibria in which beliefs are updated according to equation (15).

4.2 Solving the Model and Specification of Beliefs

We analyze the model's properties using the second-order approximation of equation (1) discussed in Section 3. That approximation works with the linearized map from period $t - 1$ beliefs to the period- t equilibrium.²²

²¹We have already embedded $Y_t = C_t + G_t$ in the construction of F . However, for goods market clearing to be consistent with optimization by final goods producers and by producers of homogeneous labor services, it must be that $\tilde{P}_{f,t}/P_t = \left(\frac{1-\xi_P\pi_t^{\epsilon-1}}{1-\xi_P}\right)^{\frac{1}{1-\epsilon}}$ and that $\tilde{W}_{i,t}/P_t = \left(\frac{w_t^{1-\epsilon}-\xi_W\pi_t^{\epsilon-1}w_{t-1}^{1-\epsilon}}{1-\xi_W}\right)^{\frac{1}{1-\epsilon}}$. These considerations provide us with two additional equations. Clearing in the bond market provides our third equation, that the demand for bonds is zero.

²²Winkler (2020) discusses a related approach to approximating a learning equilibrium.

Let $\hat{z}_t$ denote the log deviation of the variable, z_t , from the zero-inflation REE steady-state value, z .²³ In case $z = 0$, then $\hat{z}_t = z_t$.

We assume that at the time people make their time- t decisions, they believe the aggregate variables, $\hat{x}_{t+j}$, have the following law of motion:²⁴

$$\hat{x}_{t+j} = \mu + \Omega \hat{a}_{t+j-1} + \nu_{t+j}, \quad j = 0, 1, \dots \quad (16)$$

Here, μ is a 3×1 vector. People believe that the 3×1 vector, ν_t , is *i.i.d.* over time and normally distributed with mean 0 and diagonal variance-covariance matrix. In Appendix B we show that the REE has the form of equation (16) with $\mu = 0$, a specific matrix Ω , and $\nu_t \equiv 0$. The presence of $\hat{a}_{t-1}$ in equation (16) reflects that the measure of price dispersion, p_{t-1}^* , and the lagged real wage, w_{t-1} are part of the state in a period- t equilibrium. Because we work with a linear approximation, people's belief about the variance of ν_t does not affect their decisions. People's believed value of μ is given by m_{t-1} , so that m_{t-1} corresponds to the vector Θ_{t-1} . For tractability, we assume that people believe that Ω is equal to its REE value in a first-order approximation to the model. At the time people make their time- t decision, they are perfectly certain in their beliefs about μ and Ω in the law of motion of $\hat{x}_{t+j}$ for $j \geq 0$. The value of ν_t is then determined by the requirement that time- t markets clear (see Appendix B for further discussion). Note that this means people's assumption that ν_t is an independent random variable is incorrect.

After seeing time- t aggregate outcomes, people update beliefs using Bayes' rule:²⁵

$$m_t = m_{t-1} + \gamma_t (\hat{x}_t - \Omega \hat{a}_{t-1} - m_{t-1}) \quad (17)$$

where $\gamma_t = (\lambda + t)^{-1}$.²⁶

Appendix B shows that $\hat{x}_t - \Omega \hat{a}_{t-1} = H m_{t-1}$ in a first-order approximation to the period- t equilibrium. So, a second-order approximation to equation (17) is given by

$$m_t = m_{t-1} + \gamma_t (H m_{t-1} - m_{t-1}) \quad (18)$$

where $\gamma_t = \gamma_{t-1} (1 - \gamma_{t-1})$, $0 < \gamma_1 < 1$. A surprising feature of equation (18) is that only the variables m_{t-1} and γ_t (and not the variables $\hat{x}_t$ and $\hat{a}_t$) appear on the right-hand-side.

²³The model has another steady state in which $\pi_t = \beta$ (see, Benhabib et al. (2001a)). This other steady state is not stable under the learning we study in this paper. See Arifovic et al. (2018) for an approach to learning in which that steady state is stable under learning.

²⁴Equation (16) is the analog of equation (14).

²⁵People never revise their belief about Ω .

²⁶At the time that people make their decision, they believe $\hat{x}_t - \Omega \hat{a}_{t-1} = \mu + \nu_t$, and think they know the value of μ with perfect certainty (their priors are characterized by infinite precision). If that were so, they would attribute all the deviation between $\hat{x}_t - \Omega \hat{a}_{t-1}$ and m_{t-1} to ν_t and set $\gamma_t = 0$. The internal inconsistency in anticipated utility is that when people update their beliefs, they set $\gamma_t > 0$, indicating that the precision in their priors is finite.

This feature of the solution exploits our assumption that people are only uncertain about intercept terms and not slope terms in equation (16).

Note that equation (18) is of the form of equation (1) with $M(m_{t-1}, \gamma_t) = Hm_{t-1}$. So, we can apply our Proposition to this environment under Assumptions 1 and 2.

4.3 Speed of Convergence

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

$$\beta = 0.995, \varepsilon = \epsilon = 4, \xi_P = \xi_W = 0.85, \chi = 1.25, \alpha_\pi = 1.5, \alpha_Y = 0.25, G = 0.2Y, \vartheta = 10^{-6}.$$

We choose β so that the steady state real interest rate is about 2 percent at an annual rate. We set $\varepsilon = \epsilon = 4$ 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. We set $\xi_P = 0.85, \alpha_\pi = 1.5, \alpha_Y = 0.25$, which are also within the range of estimates considered in the literature. For symmetry, we set $\xi_W = \xi_P$. The value of χ is chosen so that in the target-inflation ($\pi = 1$) steady state labor supply is unity.

Table 1 reports the value of b in the benchmark model as well as in the following variants of the model: (i) wages are flexible and the labor market is competitive; (ii) prices are flexible; (iii) the coefficient α_π in the Taylor rule is increased to 2.5; (iv) the coefficient α_π in the Taylor rule is reduced 1.01; and (v) the coefficient α_Y is reduced to 0.125.

We use our asymptotic results to develop intuition for the speed of convergence in the benchmark NK model with sticky prices and wages. We exploit a result for an important special case that applies to the models that we consider. For arbitrary m_{t-1} , the evolution of beliefs is determined by the full matrix H (see equation (18)). However, if the largest real part of the eigenvalues of H , b , is real and not repeated, then $Hm_{t-1} = bm_{t-1}$, if m_{t-1} is proportional to the eigenvector, q , associated with b . Moreover, apart from special exceptions, for any initial beliefs, m_0 , the non- b eigenvalues play a vanishingly small role in the dynamics of m_t for large t and eventually m_t is very nearly proportional to q .²⁷ At that point, the period t equilibrium is in effect completely characterized by the scalar, b , and equation (18) approaches:

$$m_t = (1 + \gamma_t(b - 1))m_{t-1}, \tag{19}$$

where m_{t-1} is proportional to q . In the spirit of these asymptotic results, we consider initial conditions, m_0 , which are proportional to q . We scale m_0 so that $\|m_0\| = 1$ and normalize

²⁷By “special” we mean an m_0 that suppresses the b eigenvalues. See Section 3 for an example.

m_0 so that the value of the intercept of the inflation process is positive.

Table 1 reports the values of m_0 for each variant of the model that we consider. The vector of initial beliefs has different lengths depending on the nature of nominal rigidities. When both prices and wages are sticky, people have to form beliefs about aggregate inflation, consumption, and the real wage. When only wages or prices are sticky, people only have to form beliefs about aggregate inflation and consumption.

Table 1 reports that the value of b in the benchmark model (sticky wages and prices) is 0.84. This value implies that asymptotic convergence to REE beliefs is very slow, with convergence measured in hundreds of years. We simulate the model to find the number of periods, $\hat{T}_{2/3}$, needed to close 2/3 of the gap, $\|m_0\|$, by setting $\gamma_1 = 0.5$ and iterating on equation (19).²⁸ We provide an alternative approximation that assumes $\|m_t\|$ is asymptotically proportional to t^{b-1} (see the discussion after Proposition 1). Under this approximation, the amount of time needed to close 2/3 of an initial gap in beliefs from the REE is $T_{2/3} = (1/3)^{\frac{1}{b-1}}$. For the benchmark model, $\hat{T}_{2/3}$ and $T_{2/3}$ are 2343 and 881 quarters or, more than 500 and 200 years, respectively.

Table 1: b in the simple NK model with Calvo-style nominal rigidities

	b	$T_{2/3}$	$\hat{T}_{2/3}$	Initial beliefs		
				[	$m_{C,0}$ $m_{\pi,0}$ $m_{w,0}$	]
Sticky prices and wages	0.84	881	2343	[	-0.92, 0.38, -0.05	]
Sticky prices, flexible wages	0.26	5	6	[	-0.93, 0.37,	]
Sticky wages, flexible prices	0.81	373	980	[	-0.93, 0.37	]
Higher α_π ($\alpha_\pi = 2.5$)	0.55	12	24	[	-0.99, 0.15, -0.07	]
Lower α_π ($\alpha_\pi = 1.01$)	0.97	1.5×10^{14}	>1,000k	[	-0.02, 1.00, 0	]
Lower α_Y ($\alpha_Y = 0.125$)	0.71	42	101	[	-0.97 0.22 -0.06	]

Note: b corresponds to the largest real part of the eigenvalues of the learning model, as described in section 3; $T_{2/3}$ is the amount of time it takes to close 2/3 of the initial gap in beliefs in a simulation of the model starting with the initial beliefs listed in the table. To calculate $\hat{T}_{2/3}$, we set $\gamma_1 = 0.5$ and simulate the model using $\gamma_t = \gamma_{t-1}(1 - \gamma_{t-1})$. The latter is the second-order approximation to $\gamma_t = \frac{1}{1+t}$. We denote the initial belief about the intercept in the process for C_t by $m_{C,0}$. We define $m_{\pi,0}$ and $m_{w,0}$ similarly.

Source: Authors' calculations.

Below, we provide intuition for the forces governing the asymptotic speed of convergence to an REE. We do so by considering the impact on b of several perturbations in model parameters.

²⁸When we simulate the model, we use the second-order approximation to $\gamma_t = \frac{1}{1+t}$, which is given by $\gamma_t = \gamma_{t-1}(1 - \gamma_{t-1})$. We also simulated the model using $\gamma_t = \frac{1}{1+t}$. Our qualitative conclusions are robust to using either specification. When we simulated the model with $\gamma_t = \frac{1}{1+t}$, $\hat{T}_{2/3}$ is quantitatively a little smaller than reported in Table 1. See footnote 17 for further discussion.

Sticky Prices and Flexible Wages It is useful to begin by considering the case where prices are sticky and wages are flexible. According to the value of m_0 in Table 1, in period 1, people believe that the intercept for consumption is lower and inflation is higher than in the REE. Table 1 reports that expectations are not very self-fulfilling, $b = 0.26$, so that beliefs converge quickly to their REE values: $T_{2/3}(\hat{T}_{2/3}) = 5(6)$ quarters.

Consider the effect of high expected future inflation. Firms understand that high future inflation implies high future nominal wages, raising their desired future prices. In the presence of sticky prices, firms front-load price increases, thereby inducing *upward* pressure on period-1 inflation.

Next, consider the effect of low expected future consumption. Low expected aggregate consumption implies that households will have a high marginal utility of consumption and a high labor supply in the future. Other things equal, firms expect low nominal wages in the future. This expectation induces them to lower prices in period 1. Low expected aggregate consumption means households expect low future labor income and profits. So, consistent with firms' expectations, households increase their labor supply and lower consumption. Together, these labor-market effects exert *downward* pressure on period-1 inflation.

Because of the competing forces at work, the deviation of period 1 inflation from its REE value could be higher or lower than its expected value, $m_{\pi,0}$. Suppose, as is the case for our parameterized model and our specification of m_0 , that the actual deviation in inflation is lower. Since period 1 inflation above its REE value and $\alpha_\pi > 0$, the nominal interest rate rises above its REE value. The magnitude of that rise is moderated by the fact that $\alpha_Y > 0$ and the fact that consumption is below its REE value. On net, the period 1 equilibrium value of inflation is above its REE value but substantially below the level implied by $m_{\pi,0}$. Actual consumption is lower than its REE value, but substantially above the level implied by $m_{C,0}$.

At the end of period 1, people update their expectations based on period 1 market outcomes. The muted change in period-1 consumption and inflation, relative to initial expectations, prompts 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.

Sticky Wages and Flexible Prices Consider next the case when wages are sticky and prices are flexible. Table 1 reports that convergence occurs much more slowly than the sticky-price only case price: $b = 0.81$ and $T_{2/3}(\hat{T}_{2/3}) \approx 100$ (250) years. The intuition for this slow convergence is as follows. Initial beliefs, m_0 , are similar to their values in the sticky-price-only case. High expected inflation in subsequent periods leads unions to set

higher nominal wages in the initial period to avoid a low relative wage in the future. Absent Calvo price frictions, firms set their nominal price as a markup over the nominal marginal cost which is equal to the nominal wage rate. Because nominal marginal cost is 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. Again, there are conflicting forces acting on the realized inflation rate. But, 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 speed of learning is lower.

Sticky Wages and Sticky Wages (benchmark model) Table 1 shows that the benchmark model has a value of b that is similar to the sticky wage, flexible price model. Evidently, the dynamics of sticky wages are the dominant factor determining the large value of b and the extent to which beliefs are self-fulfilling.

Alternative Values of α_π According to Table 1, learning is faster when α_π is higher. As indicated by $m_{0,\pi}$, when $\alpha_\pi = 2.5$, initial beliefs about inflation are above their REE value. As discussed above, the positive value of $m_{\pi,0}$ leads firms to raise period 1 prices. When α_π is large, monetary policy aggressively raises the interest rate in response to higher inflation. Other things equal, this rise leads to a fall in aggregate consumption. That fall induces an increase in labor supply and, a fall in the aggregate wage. The resulting decline in marginal cost moderates the rise in inflation. All these effects are increasing in α_π . So, other things equal, a higher value of α_π makes inflation beliefs less self-fulfilling and leads to a smaller value of b .

Table 1 also reports that when $\alpha_\pi = 2.5$, initial beliefs about aggregate consumption are below their REE value. As a consequence, individual households believe their income will be lower than its REE value. So, they cut back on their spending and increase their labor supply. The latter effect leads to a fall in the wage, marginal cost, and prices. When α_π is large, monetary policy cuts the interest rate aggressively, thus moderating the fall in consumption and inflation.

The net effect of the compound experiment involving $m_{\pi,0}$ and $m_{C,0}$ is that period 1 realized inflation and consumption are closer in period 1 to their REE values when α_π is equal 2.5 rather than 1.5. Consequently, when α_π is larger people revise their expectations by more, b is smaller, and convergence to the REE is quicker.

Consistent with the previous intuition, Table 1 reports that when $\alpha_\pi = 1.01$, b is higher

than its benchmark value and convergence to the REE is extraordinarily slow.

The Impact of α_Y Finally, consider the case in which $\alpha_Y = 0.125$ so monetary policy reacts less to output than in the benchmark case. Table 1 reports that a smaller value of α_Y reduces the value of b . As a result, learning is faster, with $T_{2/3}$ and $\hat{T}_{2/3}$ falling to about 10 and 25 years, respectively. The intuition is that a lower value of α_Y acts like a higher value of α_π , which accelerates the rate of convergence. In the polar case of $\alpha_Y = 0$, b falls to 0.64, and $T_{2/3}$ and $\hat{T}_{2/3}$ fall to about 5 and 0.5 years, respectively.

5 Speed of Learning in an Empirically Plausible DSGE Model

In this section, we show that our results about the speed of convergence in the simple NK model hold in an empirically plausible, medium-sized DSGE model. The model we analyze is similar to the one in CEE and has the following features: sticky prices and wages, internal habit formation in consumption, endogenous capital utilization rates, and adjustment costs on investment. Monetary policy is governed by a Taylor rule with interest rate smoothing. In Appendix D, we provide a detailed specification of this model. We set all non-policy parameters to the point estimates reported by CEE. For convenience, these values are reported in Appendix D. For the monetary policy rule, we set the coefficients on the inflation rate and the output gap to the values we used in the benchmark simple NK model analysis. We set the interest rate smoothing parameter to 0.8, as in CEE.

We approximate the learning equilibrium using the same method as in the simple NK model. As in the latter model, we assume that people know how aggregate prices and quantities react to their lagged values. However, they must learn about the intercepts of those processes. As above, we calculate the speed of convergence of those beliefs to their REE values.

Table 2 reports the values of b , $T_{2/3}$, and $\hat{T}_{2/3}$ for nine versions of the model. The speed of convergence to the REE beliefs in the benchmark model, labeled “CEE (2005),” is very slow. Measured by $T_{2/3}$ and $\hat{T}_{2/3}$, it takes roughly 800 years to close 2/3 of the gap between the initial beliefs and the REE beliefs. So, moving from the simple NK model to a more empirically plausible model increases the time it takes to converge to the REE. The intuition for the role of sticky prices and the monetary policy parameters, α_π and α_Y , is qualitatively similar to the intuition provided in the context of the simple NK model.

As in the simple NK model, the parameters of the monetary policy rule and sticky wages

are key determinants of the speed of learning. Allowing wages to be flexible dramatically reduces $T_{2/3}$ from more than 33,000 to 22 (roughly 5 years). Table 2 reports that, conditional on the monetary policy rule, removing interest rate smoothing, habit formation, and adjustment costs on investment has little effect on the speed of convergence. Reducing α_Y to 0.125 also does not significantly reduce the speed of learning. But, in the polar case of $\alpha_Y = 0$, $b = 0.27$ and learning is rapid ($T_{2/3}$ is about one year). This extreme result depends on habit formation. If we remove this feature from the model, b is equal to 0.55, essentially the same value as in the simple NK model with $\alpha_Y = 0$ and sticky wages and prices.

Table 2: b in variants of the CEE (2005) model with Calvo-style nominal rigidities

	b	$T_{2/3}$	$\hat{T}_{2/3}$
CEE (2005)	0.89	33.1k	89.8k
Sticky prices, flexible wages	0.64	22	48
Sticky wages, flexible prices	0.89	15.6k	42.1k
No interest rate smoothing	0.88	12.2k	33.0k
No habit formation	0.90	37.3k	101k
No investment adjustment costs	0.90	44.3k	120k
Higher α_π ($\alpha_\pi = 2.5$)	0.70	41	98
Lower α_π ($\alpha_\pi = 1.01$)	0.98	2.3×10^{26}	$> 1,000k$
Lower α_Y ($\alpha_Y = 0.125$)	0.81	376	990

Note: b corresponds to the largest real part of the eigenvalues of the learning model, as described in section 3, $T_{2/3}$ is the approximate time it takes to close $2/3$ of the initial gap in beliefs using t^{b-1} as an approximation for $\|m_t\|$, $\hat{T}_{2,3}$ is the amount of time it takes to close $2/3$ of the initial gap in beliefs in a simulation of the model. To implement no sticky wages, no sticky prices, no interest rate smoothing, no habit formation, and no adjustment costs on investment, we divide the benchmark values of the relevant parameters by 100. To calculate $\hat{T}_{2/3}$, we set $\gamma_1 = 0.5$ and simulate the model using $\gamma_t = \gamma_{t-1} (1 - \gamma_{t-1})$. The latter is the second-order approximation to $\gamma_t = \frac{1}{1+t}$.
Source: Authors' calculations.

6 Robustness of the Analysis

In this section, we assess the robustness of our conclusions to various perturbations of the benchmark analysis. The key perturbations involve model economies in which the ZLB is binding. The ZLB is of particular interest because it allows us to analyze the role of the Taylor principle in governing the speed of learning. For computational reasons, we adopt Rotemberg (1982)-style rather than Calvo (1983)-style nominal rigidities. Subsection 6.1 assesses the robustness of our results for the benchmark NK model to this change. In subsection 6.2 we analyze the speed of learning in the simple NK model when the ZLB is binding. In subsection 6.2.1 we compute b using a linear approximation to the model. In subsection 6.2.2 we redo the analysis of subsection 6.2.1 in the fully nonlinear version of the simple NK model. Subsection 6.2.3 assesses the robustness of the results of section 6.2.2 to

assuming that learning is internally consistent in the sense of Adam and Marcet (2011). We find that our key conclusions about the speed of learning are robust to the perturbations we consider.

6.1 Rotemberg-Style Nominal Rigidities

Under RE, Calvo-style and Rotemberg-style nominal rigidities are equivalent to a first-order approximation (see, for example, Born and Pflaier (2020)). However, they need not be equivalent in a learning equilibrium (see Definition 3). We choose adjustment cost parameters so that the model with Rotemberg-style nominal rigidities is equivalent to a first-order approximation, in an REE, to the model with Calvo-style nominal rigidities. Wherever possible, we maintain the notation used in Sections 4 and 5.

For the simple NK model with Rotemberg-style nominal rigidities, each household is a monopoly provider of a differentiated labor input, $N_{h,t}$, that is aggregated by competitive firms using the technology $N_t = \left(\int_0^1 N_{h,t}^{\frac{\epsilon-1}{\epsilon}} dh \right)^{\frac{\epsilon}{\epsilon-1}}$. Households choose $C_{h,t}$, $b_{h,t}$, and $w_{h,t}$ to maximize

$$\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \left\{ \log(C_{h,t+j}) - \frac{\chi}{1+\zeta} N_{h,t+j}^{1+\zeta} - \frac{\Phi_W}{2} \left(\frac{w_{h,t}}{w_{h,t-1}} \pi_t - 1 \right)^2 \right\}$$

subject to their budget constraints

$$C_{h,t} + \frac{b_{h,t}}{R_t} \leq \frac{b_{h,t-1}}{\pi_t} + (1 + \tau_w) w_{h,t} N_{h,t} - \tau_t$$

and demand curves $N_{h,t} = \left(\frac{w_{h,t}}{W_t} \right)^{-\epsilon} N_t$. Because households are identical across h , they all pick the same values for $C_{h,t}$, $b_{h,t}$, and $w_{h,t}$.²⁹

The intermediate goods firms set their prices, $\tilde{P}_{f,t}$, to maximize

$$\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \Lambda_{t+j} \left\{ \left(\frac{P_{f,t+j}}{P_{t+j}} - (1 - \tau_f) \frac{W_{t+j}}{P_{t+j}} \right) \left(\frac{P_{f,t}}{P_{t+j}} \right)^{-\epsilon} Y_{t+j} - \frac{\Phi_P}{2} \left(\frac{P_{f,t}}{P_{f,t-1}} - 1 \right)^2 (C_{t+j} + G_{t+j}) \right\}.$$

Because firms are identical across f , they all pick the same value of $\tilde{P}_{f,t}$. Goods market clearing implies $Y_t = (C_t + G_t) (1 + \frac{\Phi_P}{2} (\pi_t - 1)^2)$. The rest of the model is the same as the one in section 4.1 with $p_t^* = 1$ and price adjustment costs subtracted from firm profits. Beliefs are updated as in section 4.2, with $\Theta_{t-1} = m_{t-1} = \left[m_{C,t-1}, m_{\pi,t-1}, m_{w,t-1} \right]'$.

We implement Rotemberg-style adjustment costs in the CEE model in a similar way. Table 3 reports the values of b in variants of the simple NK model and CEE model with Calvo-style nominal rigidities and in the model with Rotemberg-style nominal rigidities. The

²⁹We treat bond decisions in the same way as discussed in subsection 4.1.2.

value of b for the two different ways of modeling nominal rigidities is very similar.

Table 3: b with Calvo-style and Rotemberg-style nominal rigidities

	b (Calvo)	b (Rotemberg)
Simple NK model	0.84	0.84
Simple NK model: only sticky prices	0.26	0.33
Simple NK model: only sticky wages	0.81	0.83
CEE (2005)	0.89	0.93
CEE (2005): only sticky prices	0.64	0.74
CEE (2005): only sticky wages	0.89	0.93

Note: The table displays values for b in the two models. b corresponds to the largest real part of the eigenvalues of the learning model, as described in section 3. The values in the column labeled “ b (Calvo)” are the same as those listed in Tables 1 and 2. The values in the column labeled “ b (Rotemberg)” are from the version of the model with Rotemberg-style nominal rigidities. Source: Authors’ calculations.

6.2 The ZLB

In this subsection, we analyze the speed of learning when the ZLB is binding. Up to the assumption of RE, the environment is the same as the one considered in Eggertsson and Woodford (2003). To facilitate comparison with most of the related literature, we focus on the case of sticky prices and flexible wages, $\Phi_W = 0$. We begin by analyzing the properties of a linear approximation to the model’s solution at the ZLB.

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

$$R_t = \max \{1, \beta^{-1} + \alpha_\pi (\pi_t - 1)\}, \quad (20)$$

where $\alpha_\pi \beta > 1$, and the max operator reflects the ZLB constraint.

In period t , households and firms discount next period’s utility by $1/(1 + r_t)$. In steady state, $r_t = r^{ss} > 0$. We assume that, until period 0, the economy has been in the non-stochastic steady state REE in which the nominal interest rate is positive. At the end of period 0, unexpectedly, people know that everyone’s discount rate has dropped, $r_1 = r^\ell < r^{ss}$. People correctly understand that r_{t+1} is drawn from a two-state Markov chain:

$$\begin{aligned} \Pr [r_{t+1} = r^\ell | r_t = r^\ell] &= p, \quad \Pr [r_{t+1} = r^{ss} | r_t = r^\ell] = 1 - p, \\ \Pr [r_{t+1} = r^\ell | r_t = r^{ss}] &= 0, \end{aligned} \quad (21)$$

for $t = 1, 2, \dots$. We assume that people know that when r_t reverts to r^{ss} , the economy returns to the steady state REE.³⁰ This assumption allows us to focus attention on learning while the ZLB is binding.

³⁰There is another RE steady state in which deflation occurs and the ZLB is binding (see Benhabib et al. (2001a)). 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

We assume that, when people are confronted with the unprecedented drop in r_t , they become very uncertain about the distribution of aggregate variables. Arguably, it is reasonable to assume that the priors of their means are centered on the historical values of those variables (i.e., π^{ss} and C^{ss}). While prior variances were degenerate before the shock to r_t , they are large when people become aware of the drop in r . With this specification of initial beliefs, we in effect focus on the real-time evolution of beliefs instead of only focusing on asymptotics.

In the applications below, we use the same parameter values as in section 6.1. In addition, we specify the two new parameters, p and r^ℓ , to be $p = 0.8$ and $r^\ell = -0.0015$.³¹ In the $R > 1$ steady-state REE, $C^{ss} = 0.8$, $\pi^{ss} = 1$, $N^{ss} = 1$. We assume that the initial precision of the priors is $\lambda = 1$. Appendix E describes the details of how we solve the learning model using a second-order approximation method, which involves a first-order approximation to the solution of the nonlinear NK model.

6.2.1 Speed of Convergence Using the Second-Order Approximation Solution Method

Under RE, the ZLB equilibrium is characterized by constant values of consumption, C^ℓ , and inflation, π^ℓ . The model has multiple equilibria when $r_t = r^\ell$. Here, we focus on the unique learnable ZLB equilibrium and linearize the model around that equilibrium.³²

Our assumption about the mean of people's priors, m_0 , at the end of period 0 implies that

$$m_0 = [m_{0,C}, m_{0,\pi}] = [\log(C^{ss}/C^\ell), \log(\pi^{ss}/\pi^\ell)].$$

Our parameterized model implies that $b = 0.92$. It turns that m_0 is roughly proportional to the eigenvector associated with b .

Our key result is that the learning economy converges very slowly when the ZLB is binding: $T_{2/3} = 1.7$ million and $\hat{T}_{2/3} = 2.6$ million. 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 at the end of $t - 1$ is that inflation will be low in period t , causing firms to want to lower prices in period t . When the ZLB binds, the nominal

by Benhabib et al. (2001a). We examined the case in which the economy jumps to the Benhabib et al. (2001a) steady state when r_t returns to r^{ss} and found similar results. This is perhaps not surprising because the two steady states are relatively close to each other.

³¹It 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 slower convergence.

³²See Appendix E for a discussion of the linearization procedure. See Christiano et al. (2018) and Appendix G for further discussion on REE multiplicity.

interest rate can't adjust in the current period to prevent the fall in inflation. Consequently, expectations are more self-fulfilling, and people only adjust their beliefs very slowly.

6.2.2 Speed of Convergence in a Fully Nonlinear Model

In this subsection, we assess the robustness of our results to using a fully non-linear solution of the simple NK model. In the model, macroeconomic aggregates move by large amounts when the ZLB is binding. This fact makes the ZLB environment particularly useful for assessing the importance of nonlinearities.

As above, households and firms form beliefs about $x_t = [\log(C_t), \log(\pi_t)]'$. While $r_t = r^\ell$, people believe that

$$\begin{bmatrix} \log(C_t) \\ \log(\pi_t) \end{bmatrix} = \begin{bmatrix} \mu_C \\ \mu_\pi \end{bmatrix} + \begin{bmatrix} \nu_{C,t} \\ \nu_{\pi,t} \end{bmatrix} \quad (22)$$

where $\nu_{C,t} \sim N(0, \sigma_C^2)$, and $\nu_{\pi,t} \sim N(0, \sigma_\pi^2)$. These distributions are independent across time. The above expression is a special case of equation (14).

Our baseline assumption is that when people are confronted with an unprecedented observation, they become very uncertain about how market-determined variables will evolve. Their prior about μ_i conditional on σ_i^2 is Normal, parameterized with a mean, m_i , and variance, σ_i^2/λ_i , where λ_i 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, with shape and scale parameters, α_i and $(\psi_i^2(\alpha_i + 1/2))$, respectively (see Appendix F.2.1 for further details). The prior for σ_i^2 is not exactly an inverse-gamma distribution because we truncate the support of σ_i^2 so that $\mathbb{E}_t[C_{t+1}]$ and $\mathbb{E}_t[\pi_{t+1}]$ have finite values.³³ We find it convenient to express the scale parameter in this way because ψ_i is a consistent estimator for σ_i . The joint density of μ_i, σ_i^2 is proportional to the Normal inverse-gamma distribution. We collect the prior parameters in the vector Θ_0 .

Under anticipated utility, at time t people set the values of μ_C and μ_π equal to the elements of $m_{t-1} = [m_{C,t-1}, m_{\pi,t-1}]'$. Similarly, they set σ_C^2 and σ_π^2 to $\psi_{C,t-1}^2$ and $\psi_{\pi,t-1}^2$, respectively. In making their time- t decisions, people behave as though they know with certainty that those parameters will remain unchanged in the future.

We set the initial value of Θ_0 as follows:

$$\begin{aligned} \Theta_0 &= \left(m_{C,0}, m_{\pi,0}, 1/\lambda_C, 1/\lambda_\pi, \psi_C, \psi_\pi, \alpha_C, \alpha_\pi \right)' \\ &= \left(\log(C^{ss}), \log(\pi^{ss}), 1, 1, 0.02, 0.02, 1/2, 1/2 \right)' \end{aligned}$$


³³See Geweke (2001) for related discussion. Bakshi and Skoulakis (2010) analyzed truncation of the distribution of the variance.

These priors are diffuse in two senses. First, if ψ_C, ψ_π were the standard deviation of $\log C_t$ and $\log \pi_t$, then the distribution of $\log C_t$ and $\log \pi_t$ would have a quarterly standard deviation of two percent. Second, the assumed values $\left(1/\lambda_C \ 1/\lambda_\pi \ 1/\alpha_C \ 1/\alpha_\pi \right)$ imply low precision in people's beliefs about both the mean and standard deviation of $\log(C_t)$ and $\log(\pi_t)$, in the sense that people's beliefs are relatively responsive to new data. The posterior distribution is also proportional to the Normal inverse-gamma distribution, and the function, L , in equation (15) can be constructed using standard updating formulas (see Appendix F for a detailed description of our computational approach to solving the nonlinear model).

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 at the ZLB. Throughout the simulation, people behave as though r_t will revert to r^{ss} with probability $1 - p$.

Two key features of Figure 1a are worth noting. First, the outcomes for consumption and inflation are very different in the REE and the learning equilibria. 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 equilibrium converges very slowly to the REE. People initially change their views somewhat quickly because the small values of λ_i and α_i for $i \in \{C, \pi\}$ imply that γ_1 in equation (15) is relatively large. The rate at which they change their views slows dramatically because γ_t becomes smaller as time passes (see Figure 1b).

Figure 1: Learning equilibrium in non-linear model



Note: Sub-figure (a) shows outcomes for inflation, consumption, and the real interest rate under learning and in an REE when the natural rate of interest is low ($r_t = r^\ell < r^{ss}$). The natural rate of interest is low for the entire simulation. This is a very low probability event because the simulations cover 20 quarters. People believe that the natural rate of interest will return to its steady state level with probability $1 - p$. We use the simple NK model with only sticky prices and Rotemberg-style nominal rigidities. Sub-figure (b) shows the values of $\mathbb{E}_t \log(\pi_{t+1})$ and $\mathbb{E}_t \log(C_{t+1}/C^{ss})$ over the course of the simulation. Source: Authors' calculations.


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 expected to be low for only about 5 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. That is, our results about slow convergence when the ZLB is binding are consistent with those that we obtained with the hybrid log-linear approximation to the model's solution in section 6.2.

6.2.3 Speed of Convergence with Internalized Learning in a Fully Nonlinear Model

In this subsection, we examine the robustness of our result to the assumption that people make decisions based on internalized learning. 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 the fact that in the future they will update their views about the distribution of unknown parameters as new data are observed. Third, under anticipated utility, people ignore their uncertainty about the mean and variance of unknown parameters. In internalized learning, people observe aggregate outcomes, then update their beliefs and make their choices. In solving their problems, they fully internalize that they will update their beliefs in response to new data. See Appendix F for a description of our computational approach to solving the model under internalized learning.

Figure 2: Anticipated utility versus internalized learning



Note: The figure shows outcomes for inflation, consumption, and the real interest rate under anticipated utility and under internalized learning when the natural rate of interest is low ($r_t = r^\ell < r^{ss}$). The natural rate of interest is low for the entire simulation. This is a very low probability event because the simulations cover 20 quarters. People believe that the natural rate of interest will return to its steady state level with probability $1 - p$. We use the simple NK model with only sticky prices and Rotemberg-style nominal rigidities. Notice that the dashed black lines, labeled “Anticipated utility,” in this figure are the same as the dashed black lines in Figure 1(a).

Source: Authors’ calculations.

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 initially 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. In internalized learning, people update their beliefs before making decisions, leading them to expect lower future inflation and consumption. Because the ZLB is binding, the real interest rate is higher, leading to even lower inflation and consumption. In any event, consumption, inflation, and the real interest rate are persistently very different than the REE outcomes. So, our results on slow convergence when the ZLB is binding are robust to the use of internalized learning.

7 Does Learning Matter for Policy?

In this section, we discuss how slow learning affects the implications of monetary and fiscal policy. It is well known that the policy implications of the ZLB are dramatic in REE. 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 others 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 policy implications, we use the non-linear solution to the simple NK model when the ZLB is binding, and people use internalize learning in their decision-making. As it turns out, similar conclusions obtain if we work with anticipated utility.

7.1 The Government Purchases Multiplier


In this subsection, we compare the efficacy of fiscal policy under RE 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 while $r_t = r^\ell$. We define the multiplier in the ZLB at time $t + j$ as

$$\frac{\Delta C_{t+j} + \Delta G_{t+j}}{\Delta G_{t+j}}. \quad (23)$$

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

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 a rise in G raises inflation and expected inflation. Because R_t is constrained by the ZLB, this rise leads to a fall in the real interest rate and a rise in C_t . So, in this case, the multiplier is bigger than unity.

Figure 3: Equilibria with and without an increase in G_t



Note: Sub-figure (a) shows outcomes for inflation, consumption, and the real interest rate under internalized learning when the natural rate of interest is low ($r_t = r^\ell < r^{ss}$) and G_t is either high or at its steady state value while $r_t = r^\ell$. Sub-figure (b) shows the value of the multiplier defined in equation (23). The natural rate of interest is low for the entire simulation. This is a very low probability event because the simulations cover 20 quarters. People believe that the natural rate of interest will return to its steady state level with probability $1 - p$. We use the simple NK model with only sticky prices and

Rotemberg-style nominal rigidities.
Source: Authors' calculations.


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 falls only slightly, 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. Because learning is slow, the multiplier is close to 1 over the 20 quarters displayed.³⁴

³⁴Mertens and Ravn (2014) discuss the effects of fiscal policy in an REE that is caused by a nonfundamental shock. They also discuss learning around that equilibrium using the Euler-equation approach to learning. We discuss the equilibrium on which they focus in Appendix H.

7.2 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 $J = 1$.

Figure 4: Forward guidance under learning and in the REE



Note: The figure shows inflation, consumption, and the real interest rate paths under internalized learning when the natural rate of interest is low ($r_t = r^\ell < r^{ss}$) with and without forward guidance about R_t . Under forward guidance, the monetary authority credibly promising to set $R_t = 1$ for one period after $r_t = r^\ell$. The natural rate of interest is low for the entire simulation. This is a very low probability event because the simulations cover 20 quarters. People believe that the natural rate of interest will return to its steady state level with probability $1 - p$. We use the simple NK model with only sticky prices and Rotemberg-style nominal rigidities. Source: Authors' calculations.

Forward guidance takes the following form. In period I , r_t reverts to r^{ss} (I is stochastic) and $R_t = 1$. In all earlier periods, when $r_t = r^\ell$, people know that in period I , C_t and π_t are equal to their REE values, C^I and π^I . People also know that, in period $I + 1$, the economy reverts to the target-inflation REE steady state. In Appendix G we show that the number of REE proliferates under forward guidance. We study the learning equilibrium during the period when $r_t = r^\ell$ using the same initial values for Θ_0 as in section 6.2.2.³⁵

Figure 4 shows that the learning equilibrium under forward guidance is very similar to the learning equilibrium without forward guidance. The reason is that forward guidance has small effects on inflation and consumption in the period I (the period when r_t reverts to r^{ss}). Those small effects have little impact on the learning equilibrium in which beliefs are partially backward-looking. Clearly, there is no forward guidance puzzle under learning.³⁶ 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. Moreover, the small effects of forward guidance persist for

³⁵See Appendix F.4 for a discussion of our solution algorithm.

³⁶For a discussion of the forward guidance puzzle, see Del Negro et al. (2023).

many quarters because learning is slow.

8 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 nonstochastic, multivariate systems. We then argue that the slow convergence result emerges naturally in empirically plausible NK models, including canonical DSGE models similar to CEE (2005). When learning is slow, analyses of fiscal and monetary policies under RE can be very misleading.

We leave several promising and related topics for future research, including: connecting our results to the literature on Bayesian REE, investigating how heterogeneity may affect the speed of learning, and analyzing how survey data might be used to inform estimates of model parameters by empirically disciplining the speed of learning.

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 Jennifer Laó**, “Optimal monetary policy with informational frictions,” *Journal of Political Economy*, 2020, 128 (3), 1027–1064.
- 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.
- Bakshi, Gurdip and Georgios Skoulakis**, “Do subjective expectations explain asset pricing puzzles?,” *Journal of Financial Economics*, 2010, 98 (3), 462–477.
- Benhabib, Jess, George W Evans, and Seppo Honkapohja**, “Liquidity traps and expectation dynamics: Fiscal stimulus or fiscal austerity?,” *Journal of Economic Dynamics and Control*, 2014, 45, 220–238.
- **, Stephanie Schmitt-Grohé, and Martin Uribe**, “Monetary policy and multiple equilibria,” *American Economic Review*, 2001, 74 (2), 165–170.

- , —, and —, “The Perils of Taylor rules,” *Journal of Economic Theory*, 2001, 96 (1-2), 40–69.
- Benveniste, Albert, Michel Métivier, and Pierre Priouret**, *Adaptive algorithms and stochastic approximations*, Springer-Verlag, 1990.
- Bilbiie, Florin O**, “Neo-fisherian policies and liquidity traps,” *American Economic Journal: Macroeconomics*, 2022, 14 (4), 378–403.
- Blanchard, Olivier J**, “The wage price spiral,” *The Quarterly Journal of Economics*, 1986, 101 (3), 543–565.
- Boneva, Lena Mareen, R Anton Braun, and Yuichiro Waki**, “Some unpleasant properties of loglinearized solutions when the nominal rate is zero,” *Journal of Monetary Economics*, 2016, 84, 216–232.
- Born, Benjamin and Johannes Pfeifer**, “The new keynesian wage phillips curve: Calvo vs. rottemberg,” *Macroeconomic Dynamics*, 2020, 24 (5), 1017–1041.
- Caines, Colin and Fabian Winkler**, “Asset price beliefs and optimal monetary policy,” *Journal of Monetary Economics*, 2021, 123, 53–67.
- Calvo, Guillermo A.**, “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, 1983, 12 (3), 383–398.
- Christiano, Lawrence J., Martin S. Eichenbaum, and Charles L. Evans**, “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 2005, 113 (1), 1–45.
- , —, and **Sergio Rebelo**, “When is the Government Spending Multiplier Large?,” *Journal of Political Economy*, 2011, (February).
- Christiano, Lawrence, Martin S Eichenbaum, and Benjamin K Johannsen**, “Does the New Keynesian Model Have a Uniqueness Problem?,” Technical Report, National Bureau of Economic Research 2018.
- Christopeit, Norbert and Michael Massmann**, “Estimating structural parameters in regression models with adaptive learning,” *Econometric Theory*, 2018, 34 (1), 68–111.
- Coddington, Earl A.**, *An Introduction to Ordinary Differential Equations*, Prentice-Hall, 1961.
- Collin-Dufresne, Pierre, Michael Johannes, and Lars A Lochstoer**, “Parameter learning in general equilibrium: The asset pricing implications,” *American Economic Review*, 2016, 106 (3), 664–698.
- Eggertsson, Gauti and Michael Woodford**, “The Zero Bound on Interest Rates and Optimal Monetary Policy,” *Brookings Papers on Economic Activity*, 2003, (1).

- 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.
- Erceg, Christopher J. and Andrew T. Levin**, “Imperfect credibility and inflation persistence,” *Journal of Monetary Economics*, 2003, 50 (4), 915–944.
- Erceg, Christopher J, 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**, “Central Bank Transparency under Model Uncertainty,” *Federal Reserve Bank of New York Staff Reports*, 2005, (199).
- , “Central bank communication and the liquidity trap,” *Journal of Money, Credit and Banking*, 2010, 42 (2-3), 373–397.
- **and Bruce Preston**, “The science of monetary policy: An imperfect knowledge perspective,” *Journal of Economic Literature*, 2018, 56 (1), 3–59.
- , **Christopher Gibbs, and Bruce Preston**, “Forward guidance with unanchored expectations,” *Bank of Finland Research Discussion Paper*, 2021, (11).
- , – , **and –** , “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, Eran Guse, and Seppo Honkapohja**, “Liquidity traps, learning and stagnation,” *European Economic Review*, 2008, 52 (8), 1438–1463.
- , **Seppo Honkapohja, and Kaushik Mitra**, “Expectations, stagnation, and fiscal policy: A nonlinear analysis,” *International Economic Review*, 2022, 63 (3), 1397–1425.
- Farhi, Emmanuel and Iván Werning**, “Monetary policy, bounded rationality, and incomplete markets,” *American Economic Review*, 2019, 109 (11), 3887–3928.
- Farmer, Leland E, Emi Nakamura, and Jón Steinsson**, “Learning about the long run,” *Journal of Political Economy*, 2024, 132 (10), 3334–3377.

- Ferrero, Giuseppe**, “Monetary policy, learning and the speed of convergence,” *Journal of Economic Dynamics and Control*, 2007, 31 (9), 3006–3041.
- Geweke, John**, “A note on some limitations of CRRA utility,” *Economics letters*, 2001, 71 (3), 341–345.
- 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.
- Judd, Kenneth L**, “Approximation, Perturbation, and Projection Methods in Economic Analysis,” *Handbook of Computational Economics*, 1996.
- 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.
- Lucas, Robert E.**, “Asset Prices in an Exchange Economy,” *Econometrica*, 1978, 46 (6), 1429–1445.
- Marcet, Albert and Juan P Nicolini**, “Recurrent hyperinflations and learning,” *American Economic Review*, 2003, 93 (5), 1476–1498.
- **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.

- McCallum, Bennett T.**, “On non-uniqueness in rational expectations models: An attempt at perspective,” *Journal of monetary Economics*, 1983, 11 (2), 139–168.
- McCallum, Bennett T.**, “Inflation targeting and the liquidity trap,” in Loayza, N., and Soto, R., ed., *Inflation Targeting: Design, Performance, Challenges*, Central Bank of Chile, 2002, pp. 395–437.
- Mertens, Karel RSM and Morten O Ravn**, “Fiscal policy in an expectations-driven liquidity trap,” *The Review of Economic Studies*, 2014, pp. 1637–1667.
- Murphy, Kevin P.**, “Conjugate Bayesian Analysis of the Gaussian Distribution,” *mimeo*, 2007.
- 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.
- Sims, Christopher A.**, “Solving Linear Rational Expectations Models,” *Computational Economics*, 2001, 20, 1–20.
- Smets, Frank and Rafael Wouters**, “Shocks and frictions in US business cycles: A Bayesian DSGE approach,” *American economic review*, 2007, 97 (3), 586–606.
- Vives, Xavier**, “How fast do rational agents learn?,” *The Review of Economic Studies*, 1993, 60 (2), 329–347.
- Winkler, Fabian**, “The role of learning for asset prices and business cycles,” *Journal of Monetary Economics*, 2020, 114, 42–58.
- 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.
- Yun, Tack**, “Nominal price rigidity, money supply endogeneity, and business cycles,” *Journal of monetary Economics*, 1996, 37 (2), 345–370.
- Zellner, Arnold**, *An Introduction to Bayesian Inference in Econometrics* 1971.

A Proof of Proposition 1

Proposition 1 is an immediate result of two lemmas that we prove in this appendix. We assume that θ_t evolves according to equation (1), and we maintain Assumptions 1 and 2 in the text. Recall that Assumption 1 implies that the largest real part of the eigenvalues of $D_1M(0,0)$, which we denote by b , is strictly less than unity. We also state the lemmas in terms of any neighborhood of the origin, U , that has the property that every $\theta_0 \in U$ generates a sequence $\|\theta_t\|$ that converges to zero. Proposition 1 in Evans and Honkapohja (2000) establishes that such a U exists. Our first lemma establishes that for every $\theta_0 \in U$ the implied sequence, θ_t , satisfies part (1) of Definition 1. Our second lemma establishes that there exist uncountably many $\theta_0 \in U$ that imply a sequence, θ_t , that satisfies part (2) of Definition 1. The general strategy of the proofs of Lemmas 1 and 2 is similar to the strategy of the proofs of Propositions 1 and 2, respectively, in Evans and Honkapohja (2000).

A.1 Lemma 1

Lemma 1. *For any $0 < \delta$, if $\theta_0 \in U$ then $\lim_{t \rightarrow \infty} \frac{\|\theta_t\|}{t^{b-1+\delta}} = 0$.*

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 $J = D_1M(0,0) - I$ and $r_t = M(\theta_{t-1}, \gamma_t) - D_1M(0,0) \theta_{t-1}$. 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 $1 + \gamma_t(b - 1 + \delta/4)$. Note that

$$|1 + \gamma_t \lambda| = \sqrt{1 + 2\gamma_t(\lambda_r + \delta/4) - 2\gamma_t\delta/4 + \gamma_t^2\lambda_r^2 + \gamma_t^2\lambda_c^2}$$

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 , $-2\gamma_t\delta/4 + \gamma_t^2\lambda_r^2 + \gamma_t^2\lambda_c^2 < 0 < \gamma_t^2(b - 1 + \delta/4)^2$. So, for large enough t ,

$$|1 + \gamma_t \lambda| < 1 + \gamma_t(b - 1 + \delta/4). \quad (24)$$

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

$$\|I + \gamma_t J\| \leq 1 + \gamma_t \left(b - 1 + \frac{1}{2}\delta \right), \quad (25)$$

where $\|\cdot\|$ is the norm and we have used the inequality in 24.

Now consider r_t , which is a function of θ_{t-1} and γ_t . As in Evans and Honkapohja (1995), the maintained assumptions imply local Lipschitz continuity of r_t . Following Evans and Honkapohja (1995), who use results from Coddington (1961), for large enough t (small enough γ_t) there exists a neighborhood of the origin U_2 , where $U_2 \subseteq U_1$, such that for $\theta_{t-1} \in U_2$ we have $\|r_t\| \leq \frac{\delta}{4} \|\theta_{t-1}\|$ and

$$\|\theta_t\| \leq \|I + \gamma_t J\| \|\theta_{t-1}\| + \gamma_t \|r_t\| \leq \left(1 + \gamma_t \left(b - 1 + \frac{3}{4}\delta\right)\right) \|\theta_{t-1}\|. \quad (26)$$

The first inequality in 26 follows from sub-additivity and sub-multiplicativity of the matrix-vector norm. The second inequality follows from the inequalities in 25 and the inequality before the offset equation.

Consider $\theta_0 \in U_1$. For large enough τ_1 , $\theta_{t-1} \in U_2$ for $t \geq \tau_1$ because θ_t converges to 0. For $t - 1 \geq \tau_1$, define $\tilde{\theta}_{t-1} = \frac{\theta_{t-1}}{(t-1)^{b-1+\delta}}$. For large enough t there exists a finite, positive constant g so that

$$\begin{aligned} \|\tilde{\theta}_t\| &\leq \left(1 + \frac{1}{t-1}\right)^{1-b-\delta} \left(1 + \gamma_t \left(b - 1 + \frac{3}{4}\delta\right)\right) \|\tilde{\theta}_{t-1}\| \\ &\leq \left(1 + \frac{1-b-\delta}{t-1} + g \left(\frac{1}{t-1}\right)^2\right) \left(1 + \gamma_t \left(b - 1 + \frac{3}{4}\delta\right)\right) \|\tilde{\theta}_{t-1}\| \end{aligned} \quad (27)$$

The first line uses the inequality in 26 and follows from absolute homogeneity of the vector 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 in 27 can be written as $\|\tilde{\theta}_t\| \leq \tilde{k}_t \|\tilde{\theta}_{t-1}\|$ where $\tilde{k}_t = 1 + \frac{1-b-\delta}{t-1} + \gamma_t (b - 1 + \frac{3}{4}\delta) + \tilde{q}_t$, and $\tilde{q}_t = \frac{1-b-\delta}{t-1} \gamma_t (b - 1 + \frac{3}{4}\delta) + g \left(\frac{1}{t-1}\right)^2 (1 + \gamma_t (b - 1 + \frac{3}{4}\delta))$. For t large enough, $\tilde{k}_t < 1$. For all $\|\tilde{\theta}_{t-1}\| \neq 0$

$$\|\tilde{\theta}_t\| - \|\tilde{\theta}_{t-1}\| \leq (\tilde{k}_t - 1) \|\tilde{\theta}_{t-1}\|. \quad (28)$$

So, there exists a $\tau_2 > \tau_1$ so that if $t > \tau_2$ then $\|\tilde{\theta}_t\|$ is either strictly decreasing or equal to zero. Furthermore, for $t \geq \tau_2$ if $\|\tilde{\theta}_{t-1}\| = 0$ then $\|\tilde{\theta}_{t+s}\| = 0$ for all s . So, if we have a sequence θ_{t-1} that equals zero for any $t \geq \tau_2$, then $\|\tilde{\theta}_{t-1}\|$ converges to zero.

We now establish convergence of $\|\tilde{\theta}_t\|$ assuming that $\theta_{t-1} \neq 0$ for all $t \geq \tau_2$. Following

the strategy in Evans and Honkapohja (2000), we iterate the inequality in (28) to obtain

$$1 - \frac{\left\| \tilde{\theta}_{\tau_2-1} \right\|}{\left\| \tilde{\theta}_{\tau_2+s} \right\|} \leq \sum_{i=\tau_2}^{\tau_2+s} \left(\tilde{k}_i - 1 \right).$$

Note that $\lim_{s \rightarrow \infty} \sum_{i=\tau_2}^{\tau_2+s} |\tilde{q}_t| < \infty$. 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, $\lim_{s \rightarrow \infty} \sum_{i=\tau_2}^{\tau_2+s} \left(\tilde{k}_i - 1 \right) = -\infty$, meaning $\lim_{t \rightarrow \infty} \left\| \tilde{\theta}_t \right\| = 0$. $\square$

A.2 Lemma 2

Lemma 2. *For any $0 < \delta$, there exist uncountably many $\theta_0 \in U$ so that $\lim_{t \rightarrow \infty} \frac{\|\theta_t\|}{t^{b-1-\delta}} = \infty$.*

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 $r_t = M(\theta_{t-1}, \gamma_t) - D_1 M(0, 0) \theta_{t-1}$. Use the real and imaginary parts of the generalized eigenvectors of J to form a basis for $\mathbb{R}^k$. Define E_A to be the generalized eigenspace associated with eigenvalues with real parts less than $b-1$ and E_B to be the generalized eigenspace associated with the eigenvalues with real parts equal to $b-1$.³⁷ 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 $\theta_t = x_t + y_t$ and $r_t = r_{A,t} + r_{B,t}$. 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 exist bases so that the corresponding inner product (Euclidean) norms on E_A and E_B satisfy

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

Here, $\langle \cdot, \cdot \rangle$ is the inner product and $\|\cdot\|$ is the norm. Choose $0 < \delta_a$ and $0 < \delta_b$ so that

³⁷To ease exposition, we assume that there are eigenvalues 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 $b-1$.

$a + \delta_a < b - \delta_b$, $b - 1 + 2\delta_b < 1$, and $2\delta_b < \delta$, where δ is given in the statement of the proposition. As in Evans and Honkapohja (2000), define $\|\theta\| = \sqrt{\|x\|^2 + \|y\|^2}$.

Define $B_{\kappa_1}(0) = \{\theta : \|\theta\| < \kappa_1\}$ to be an open ball around the origin. From Proposition 1 in Evans and Honkapohja (2000), there exists $\kappa_1 > 0$ so that if $\theta_0 \in B_{\kappa_1}(0)$, then θ_t converges. We will only consider values of $\theta_0 \in B_{\kappa_1}(0)$.

As in Evans and Honkapohja (2000) for any $\epsilon > 0$, there exists a $\kappa_r < \kappa_1$ so that if $\theta_{t-1} \in B_{\kappa_r}(0)$ then $\|r_t\| \leq \epsilon \|\theta_{t-1}\|$. Also as in Evans and Honkapohja (2000), for $\beta > 0$ define

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

For $\theta_t \in C_\beta \cap B_{\kappa_r}(0)$, the Cauchy-Schwartz inequality implies

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

Select β large enough and ϵ small enough so that

$$(b-1-2\delta_b) \|\theta_t\|^2 \leq \langle Ax_t + By_t + r_{t+1}, x_t + y_t \rangle \leq (b-1+2\delta_b) \|\theta_t\|^2. \quad (29)$$

Following the strategy in Evans and Honkapohja (2000), we now show that for large enough t , and for any $\kappa \leq \kappa_r$, if $\theta_{t-1} \in C_\beta \cap B_\kappa(0)$ then $\theta_t \in C_\beta$. First, we establish that there is a τ_1 such that if $t \geq \tau_1$ and $\theta_t \in C_\beta \cap B_\kappa(0)$, then $\theta_{t+1} \in B_\kappa(0)$. Assume $\theta_t \in C_\beta \cap B_\kappa(0)$. Then the inequality in 29 implies

$$\|\theta_{t+1}\|^2 \leq (1 + \gamma_{t+1}2(b-1+2\delta_b)) \|\theta_t\|^2 + \gamma_{t+1}^2 \|Ax_t + By_t + r_{t+1}\|^2. \quad (30)$$

Note that

$$\gamma_{t+1}^2 \|Ax_t + By_t + r_{t+1}\|^2 \leq \gamma_{t+1}^2 (k_A + k_B + \epsilon)^2 \|\theta_t\|^2, \quad (31)$$

where k_A and k_B are constants such that $\|Ax_t\|^2 \leq k_A^2 \|x_t\|^2$ for any $x_t \in E_A$ and $\|By_t\|^2 \leq k_B^2 \|y_t\|^2$ for any $y_t \in E_B$. Then the inequalities in 30 and 31 imply that

$$\|\theta_{t+1}\|^2 \leq (1 + \gamma_{t+1}2(b-1+2\delta_b) + \gamma_{t+1}^2(\kappa_A + \kappa_B + \epsilon)) \|\theta_t\|^2.$$

There exists a $\tau_1 \geq 1$, which does not depend on the value of $\kappa \leq \kappa_r$, so that if $t \geq \tau_1$ then $(1 + \gamma_{t+1}2(b-1+2\delta_b) + \gamma_{t+1}^2(\kappa_A + \kappa_B + \epsilon)) < 1$. Then for $t \geq \tau_1$ we have that $\|\theta_{t+1}\|^2 \leq \|\theta_t\|^2$, which establishes that if $t \geq \tau_1$ then $\theta_{t+1} \in B_\kappa(0)$.

Now we show that there exists a $\tau_2 \geq \tau_1$ and a $\kappa \leq \kappa_r$ so that if $t \geq \tau_2$ and $\theta_t \in C_\beta \cap B_\kappa(0)$ then $\theta_{t+1} \in C_\beta \cap B_\kappa(0)$. We have already established that $\theta_{t+1} \in B_\kappa(0)$, so we focus on

showing that $\theta_{t+1} \in C_\beta$. To do this, we follow the steps in the proof of Proposition 2 in Evans and Honkapohja (2000) related to g_β .³⁸ Assume $\theta_t \in C_\beta \cap B_\kappa(0)$. Define $g_\beta(\theta_t) = \langle y_t, y_t \rangle - \beta^2 \langle x_t, x_t \rangle$. Notice that because $\theta_t \in C_\beta$ we have that $g_\beta(\theta_t) \geq 0$. We aim to show that there exists a τ_2 such that if $t \geq \tau_2$ and $g_\beta(\theta_t) \geq 0$, then $g_\beta(\theta_{t+1}) \geq 0$. Note that

$$\begin{aligned} g_\beta(\theta_{t+1}) - g_\beta(\theta_t) &= 2(\langle y_t, \gamma_{t+1}(By_t + r_{B,t+1}) \rangle - \beta^2 \langle x_t, \gamma_{t+1}(Ax_t + r_{A,t+1}) \rangle) \\ &\quad + \langle \gamma_{t+1}(By_t + r_{B,t+1}), \gamma_{t+1}(By_t + r_{B,t+1}) \rangle \\ &\quad - \beta^2 \langle \gamma_{t+1}(Ax_t + r_{A,t+1}), \gamma_{t+1}(Ax_t + r_{A,t+1}) \rangle \end{aligned}$$

From the inequalities above

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

Then we can pick $\kappa \leq \kappa_r$ so that $\langle y_t, \gamma_{t+1}(By_t + r_{B,t+1}) \rangle - \beta^2 \langle x_t, \gamma_{t+1}(Ax_t + r_{A,t+1}) \rangle \geq \gamma_{t+1} \frac{C_1}{2} \|\theta_t\|^2$, where $C_1 > 0$. Also from the inequalities above,

$$\begin{aligned} &\langle \gamma_{t+1}(By_t + r_{B,t+1}), \gamma_{t+1}(By_t + r_{B,t+1}) \rangle - \beta^2 \langle \gamma_{t+1}(Ax_t + r_{A,t+1}), \gamma_{t+1}(Ax_t + r_{A,t+1}) \rangle \\ &= \gamma_{t+1}^2 [\|By_t\|^2 + 2 \langle By_t, r_{B,t+1} \rangle + \|r_{B,t+1}\|^2 - \beta^2 (\|Ax_t\|^2 + 2 \langle Ax_t, r_{A,t+1} \rangle + \|r_{A,t+1}\|^2)] \\ &\geq \gamma_{t+1}^2 [-2 \|By_t\| \|r_{B,t+1}\| - \beta^2 (\|Ax_t\|^2 + 2 \|Ax_t\| \|r_{A,t+1}\| + \|r_{t+1}\|^2)] \\ &\geq \gamma_{t+1}^2 [-2k_B\epsilon - \beta^2 (k_A^2 + 2k_A\epsilon + \epsilon^2)] \|\theta_t\|^2 \end{aligned}$$

Then, for $C_2 = 2k_B\epsilon + \beta^2 (k_A^2 + 2k_A\epsilon + \epsilon^2) > 0$, we have $g_\beta(\theta_{t+1}) - g_\beta(\theta_t) \geq \gamma_{t+1} (C_1 - \gamma_{t+1} C_2) \|\theta_t\|^2$. For t large enough $g_\beta(\theta_{t+1}) - g_\beta(\theta_t) \geq 0$, meaning that if $g_\beta(\theta_t) \geq 0$, then $g_\beta(\theta_{t+1}) \geq 0$. So, we can pick the values of τ_2 and κ so that if $t \geq \tau_2$ and $\theta_t \in C_\beta \cap B_\kappa(0)$ then $\theta_{t+1} \in C_\beta \cap B_\kappa(0)$.

Pick a value of $\tau_3 \geq \tau_2$ so that $1 + \gamma_{\tau_3} 2(b-1-2\delta_b) > 0$. Consider a value of $\theta_{\tau_3} \in C_\beta \cap B_\kappa(0)$ with $\|y_{\tau_3}\| \neq 0$. For $t \geq \tau_3$, the first inequality in 29 implies that

$$\|\theta_{t+1}\|^2 \geq (1 + \gamma_{t+1} 2(b-1-2\delta_b)) \|\theta_t\|^2. \quad (32)$$

Notice that this inequality means that $\|\theta_t\| > 0$ for all $t \geq \tau_3$. Now consider $\tilde{\theta}_t = \frac{\theta_t}{t^{b-1-\delta}}$. For $t \geq \tau_3$, the inequality in 32, the absolute homogeneity of the inner-product norm, and

³⁸To 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 notation. Also, the roles of x_t and y_t are reversed in our notation relative to Evans and Honkapohja (2000).

Taylor's theorem imply that there exists a $g > 0$ so that

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

Because $2\delta_b < \delta$, if t is large enough then there exists a $\tilde{\omega} > 0$ so that $\left\|\tilde{\theta}_{t+1}\right\|^2 > (1 + \frac{\tilde{\omega}}{t}) \left\|\tilde{\theta}_t\right\|^2$. So, it must be that $\lim_{j \rightarrow \infty} \left\|\tilde{\theta}_j\right\| = \infty$.

We need to establish that any neighborhood of the origin contains uncountably many values of θ_0 that generates a sequence for θ_t so that $\theta_{\tau_3} \in C_\beta \cap B_\kappa(0)$ with $\|y_{\tau_3}\| \neq 0$. The argument involving the inverse function theorem given at the end of the proof of Proposition 2 in Evans and Honkapohja (2000) establishes this result. $\square$

B 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 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 $\hat{x}_t = \log(x_t/x)$. Otherwise, we define $\hat{x}_t = x_t$.

For a given model, we assume that there is a bounded linear REE in which

$$\hat{a}_{t+j} = \Theta_{1,RE}\hat{a}_{t+j-1} + \Theta_{\epsilon,RE}\epsilon_{t+j} \quad (33)$$

for all $j \geq 0$. Here, $\hat{a}_{t+j}$ is a vector (length n_a) of aggregate variables in (log-)deviation from steady state and ϵ_{t+j} is a vector (length n_ϵ , possibly empty) of standard Normal structural shocks. In equation (33), 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, people have beliefs about a vector (length n_x) of aggregate variables, $x_{t+j} \subseteq a_{t+j}$. We assume that those beliefs are of the form

$$\hat{x}_{t+j} = c_x + \Gamma_{1,x}\hat{a}_{t+j-1} + \Psi_{x,\epsilon}\epsilon_{t+j} + \Psi_{x,\nu}\nu_{t+j}. \quad (34)$$

People are uncertain about c_x and $\Psi_{x,\nu}$. We assume that people know $\Gamma_{1,x}$ and $\Psi_{x,\epsilon}$ to be their REE values, consistent with equation (33). Agents also think that there is another vector (length n_x) of shocks, ν_{t+j} , that could affect $\hat{x}_{t+j}$. Agents believe that ν_{t+j} 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+j} only affect the associated element of $\hat{x}_{t+j}$. In an REE each element c_x and $\Psi_{x,\nu}$ is zero. Note that the elements of ν_{t+j} do not appear in the dynamics of $\hat{a}_{t+j}$ in equation (33).

Under the anticipated utility assumption, agents believe with certainty that

$$\hat{x}_{t+j} = c_{x,t-1} + \Gamma_{1,x}\hat{a}_{t+j-1} + \Psi_{x,\epsilon}\epsilon_{t+j} + \Psi_{x,\nu,t-1}\nu_{t+j}. \quad (35)$$

for some vector $c_{x,t-1}$ (length n_x) and matrix $\Psi_{x,\nu,t-1}$ (size $n_x \times n_x$). Additionally, we assume that agents know $n_a - n_x$ relationships that allow them to uniquely determine a_{t+j} given x_{t+j} through the relationship $a_{t+j} = F(x_{t+j}, \epsilon_{t+j}, a_{t+j-1})$, and that these relationships must hold in each period. Here, F embeds the associated identities for the elements of x_{t+j} . The (log-)linear version of this mapping is given by

$$\hat{a}_{t+j} = \hat{F} \begin{bmatrix} \hat{x}'_{t+j}, & \epsilon'_{t+j}, & \hat{a}'_{t+j-1} \end{bmatrix}', \quad (36)$$

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. We gather the agent-specific variables into the vector h_{t+j} (length n_h). We specify h_{t+j} so that it includes the versions of each element of x_{t+j} that would be implied by the agent-specific choices. An example of an element of h_{t+j} corresponding to an aggregate quantity in a_{t+j} in the simple NK model is household-specific consumption (in h_{t+j}) corresponding to aggregate consumption (in x_{t+j}).

After (log-)linearization, we assume agent-specific optimality conditions and budget constraints can be written as n_h equations of the form

$$\Gamma_{0,h}\hat{h}_{t+j} + \Gamma_{0,e}\hat{e}_{t+j} + \Gamma_{0,a}\hat{a}_{t+j} = \Gamma_{1,h}\hat{h}_{t+j-1} + \Gamma_{1,ha}\hat{a}_{t+j-1} + \Gamma_{1,h\epsilon}\epsilon_{t+j}. \quad (37)$$

Here, the vector $\hat{e}_{t+j} = [\mathbb{E}_{t+j}\hat{h}'_{t+j+1}, \mathbb{E}_{t+j}\hat{a}'_{t+j+1}]'$ and $\mathbb{E}_{t+j}$ is the expectation operator conditional on agents' beliefs and the information they will have at time $t+j$. Those beliefs are given by equation (35). Note that we assume that the model does not contain any agent-specific shocks in the sense that all decisions makers are always solving the same problem (though they do not know that). We emphasize that all of the matrices in equation (37) take the values they would in an REE. Deviations from rational expectations only appear through the coefficients in equation (35). Defining the vector $\eta_{t+j} = [(\hat{h}_{t+j} - \mathbb{E}_{t+j-1}\hat{h}_{t+j})', (\hat{a}_{t+j} - \mathbb{E}_{t+j-1}\hat{a}_{t+j})']'$, we also need to impose the identity

$$\eta_{t+j} = \begin{bmatrix} \hat{h}_{t+j}, & \hat{a}_{t+j} \end{bmatrix}' - \hat{e}_{t+j-1} \quad (38)$$

Letting $\hat{y}_{t+j} = [\hat{h}'_{t+j}, \hat{a}'_{t+j}, \hat{e}'_{t+j}]'$ (which has length $2(n_h + n_a)$), the $2(n_h + n_a)$ equations

in (35)-(38) can be written as

$$\Gamma_0 \hat{y}_{t+j} = \Gamma_1 \hat{y}_{t+j-1} + C c_{x,t-1} + \Psi_\epsilon \epsilon_{t+j} + \Psi_\nu \Psi_{x,\nu,t-1} \nu_{t+j} + \Pi \eta_{t+j}. \quad (39)$$

where the matrices without subscripts are independent of the values of $c_{x,t-1}$ and $\Psi_{x,\nu,t-1}$ and are equal to their values in an REE. Letting $z_{t+j} = [\epsilon'_{t+j}, \nu'_{t+j}]'$, this is exactly the form of the system studied by Sims (2001). The solution to this system described in Sims (2001) ensures that $\mathbb{E}_{t+j} \eta_{t+j+1} = 0$. The analysis in Sims (2001) is focused on RE models. It is worth stressing that the methodology we outline here could be used to compute decision rules for h_{t+j} in an REE by setting $c_{x,t-1}$ and $\Psi_{x,\nu,t-1}$ to be their REE values.

Under the conditions discussed in Sims (2001), there is a unique solution to the system of equations in (39) that can be expressed as

$$\hat{y}_{t+j} = \Theta_1 \hat{y}_{t+j-1} + \Theta_C c_{x,t-1} + \Theta_\epsilon \epsilon_{t+j} + \Theta_\nu \Psi_\nu \Psi_{x,\nu,t-1} \nu_{t+j}. \quad (40)$$

This system represents how people believe they will act conditional on their beliefs. The time $j = 0$ version of this equation represents how they will act in period t .

Importantly, equation (40) contains the mapping

$$\hat{h}_t = \Theta_{1,h} \hat{y}_{t-1} + \Theta_{C,h} c_{x,t-1} + \Theta_{h,\epsilon} \epsilon_t + \Theta_{h,\nu} \Psi_\nu \Psi_{x,\nu,t-1} \nu_t \quad (41)$$

that defines decision rules for $\hat{h}_t$ that are optimal from the perspective of the agents in the model, given $\hat{y}_{t-1}$, 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 (39), are equal to their values used when computing an REE it must be that $\Theta_{1,h}$, $\Theta_{C,h}$, and $\Theta_{h,\epsilon}$ are also equal to their values in an REE. So, we do not include a time subscript on these matrices.³⁹ Unlike in an REE, the system contains perceived shocks, ν_t . In the period equilibrium, the values of ν_t are determined by the requirement that markets clear, 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 (41) as

$$\hat{h}_t = \Theta_{1,ha} \hat{a}_{t-1} + \Theta_{C,h} c_{x,t-1} + \Theta_{h,\epsilon} \epsilon_t + \Theta_{h,\nu} \Psi_\nu \Psi_{x,\nu,t-1} \nu_t. \quad (42)$$

The vector $\hat{e}_{t-1}$ does not appear on the right-hand side of equation (42) because in the solution to the linear model its terms are a linear combination of ν_t and ϵ_t .

³⁹To 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 follows from equation (45) of Sims (2001) that Θ_1 and Θ_C in equation (40) are equal to their REE value. Additionally, if we define $\Psi = [\Psi_\epsilon, \Psi_\nu \Psi_{x,\nu,t-1}]$, then we can use equation (45) of Sims (2001) (and the notation therein) to conclude that Θ_ϵ and Θ_ν are equal to their values in an REE.

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

$$\mathcal{R}\hat{h}_t = \hat{x}_t. \quad (43)$$

Here, $\mathcal{R}$ is a matrix of ones and zeros. In a period equilibrium, equation (43) will hold. From the perspective of the agents, we can use equations (35), (42), and (43) to obtain

$$(\mathcal{R}\Theta_{h,\nu}\Psi_\nu - I)\Psi_{x,\nu,t-1}\nu_t = c_{x,t-1} - \mathcal{R}\Theta_{C,h}c_{x,t-1}. \quad (44)$$

If $(\mathcal{R}\Theta_{h,\nu}\Psi_\nu - I)$ is invertible, then the equilibrium values of $\Psi_{x,\nu,t-1}\nu_t$ can be determined uniquely. In practice we have not encountered an example in which invertibility fails. The values $\Psi_{x,\nu,t-1}\nu_t$ induce agents to make decisions that lead to market clearing.

In a learning equilibrium, after agents observe $\hat{a}_t$ and ϵ_t and make their decisions in the period equilibrium, they update their beliefs using Bayes rule. Recall that agents are only learning about c_x and $\Psi_{x,\nu}$ in equation (34) and that the elements of ν_t are perceived to be independent of one another and over time. We assume that agents beliefs are proportional to 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 precision in these priors.⁴⁰ Then

$$c_{x,t} = c_{x,t-1} + \gamma_t(\hat{x}_t - \Gamma_1\hat{a}_{t-1} - \Psi_{x,\epsilon}\epsilon_t - c_{x,t-1}) = c_{x,t-1} + \gamma_t\Psi_{x,\nu,t-1}\nu_t. \quad (45)$$

Here, $\gamma_t = 1/(\lambda_0 + t)$ where λ_0 determines the precision of the time 1 prior on each element of c_x . From equation (44), the values of $\Psi_{x,\nu,t-1}\nu_t$ are do not depend on $\hat{a}_{t-1}$, ϵ_t , or $\Psi_{x,\nu,t-1}$. Therefore, in a learning equilibrium in which $\Gamma_{1,x}$ and $\Psi_{x,\epsilon}$ are believed with certainty to be their REE values, the mean of agents' beliefs about c_x evolves in a non-stochastic manner and $c_{x,t}$ does not depend on $\hat{a}_{t-1}$. We can write the evolution of $c_{x,t}$ as

$$c_{x,t} = c_{x,t-1} + \gamma_t(H - I)c_{x,t-1}. \quad (46)$$

Equation (46) is in the form of equation (1) in the main text. Because $M(\cdot, \cdot)$ is linear in $c_{x,t-1}$ the smoothness assumptions stated in the text related to $M(\cdot, \cdot)$ hold.

⁴⁰The priors are *proportional* to the Normal-inverse-gamma priors because we assume that agents only contemplate variances for the elements of ν_t up to some finite upper limit. This ensures that expectations of the levels of variables remain finite. At the expense of additional notation, the assumption about the same precision across the priors can be relaxed. This change implies that M in equation (1) would take both $c_{x,t-1}$ and γ_t as arguments, and has no effect on the asymptotic rate of convergence of $c_{x,t}$.

C 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 widely used.⁴¹ The Euler-equation approach to learning stands in contrast to the anticipated utility approach of Preston (2005) in which households and firms solve their infinite horizon problems given beliefs.

C.1 The Linear Approximation to the Equilibrium in the Simple NK Model with Only Sticky Prices Under Euler-Equation Learning

Consider the simple NK model with only sticky prices. The aggregate Euler equations from the REE are given by

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

These equations hold under either Calvo-style or Rotemberg-style adjustment costs so long and Φ is chosen accordingly and we do not treat p_t^* as a state variable in the Calvo model because it is always equal to 1 in a first-order approximation around $\pi = 1$. Under the assumption that people are learning 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

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

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

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

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 $b = 0.26$ and $T_{2/3} = 5$. We conclude that modeling agents infinite-horizon decision problems, rather than using aggregate Euler equations from the

⁴¹An earlier version of our paper—Christiano et al. (2018)—used a version of the Euler-equation approach to learning.

REE, can have important implications for conclusions regarding the speed of learning.

C.2 Euler-Equation Learning and the ZLB

Consider the simple NK model with only sticky prices and Rotemberg-style nominal rigidities when the ZLB is binding (as in Section 6.2). Let $\tilde{X}_t$ denote the value of X_t in period t when $r_t = r^\ell$ and the ZLB is binding. The aggregate version of the household intertemporal Euler equation is

$$\frac{1}{\tilde{C}_t} = \frac{1}{1 + r^\ell} \mathbb{E}_t \frac{p}{\tilde{\pi}_{t+1} \tilde{C}_{t+1}} + \frac{1}{1 + r^\ell} \mathbb{E}_t \frac{1 - p}{\pi^{ss} C^{ss}}.$$

The aggregate version of the firm optimality condition is

$$\begin{aligned} (\tilde{\pi}_t - 1) \tilde{\pi}_t = & \frac{\varepsilon - 1}{\Phi_P} \left[\chi \tilde{C}_t (\tilde{C}_t + \tilde{G}_t) \left(1 + \frac{\Phi_P}{2} (\tilde{\pi}_t - 1)^2 \right) - 1 \right] \left(1 + \frac{\Phi_P}{2} (\tilde{\pi}_t - 1)^2 \right) \\ & + \frac{p}{1 + r^\ell} \mathbb{E}_t (\tilde{\pi}_{t+1} - 1) \tilde{\pi}_{t+1} \left(\frac{\tilde{C}_{t+1} + \tilde{G}_{t+1}}{\tilde{C}_t + \tilde{G}_t} \right) \frac{\tilde{C}_t}{\tilde{C}_{t+1}}. \end{aligned}$$

The REE at the ZLB has no state variables. So, we can linearize these equations around the values of aggregate variables taken in the REE and use the Euler equation approach to learning. When we compute the value of b under the Euler equation learning approach, we get $b = 0.98$. This is a larger value for b than under anticipated utility. So, learning is slow using either approach when we have decreasing gain as in equation (1).

D The CEE (2005) Model with Calvo-Style Nominal Rigidities

In this appendix we detail the version of the CEE (2005) model that we use in the main text. The version of the model in this appendix has Calvo-style nominal rigidities. The simple NK model and the models with Rotemberg-style nominal rigidities involve minor modifications to these equations and use the parameter values discussed in the main text.

D.1 Households

There is a continuum of households indexed by $h \in (0, 1)$ and a continuum of household members indexed by $i \in (0, 1)$. There is perfect consumption insurance within the household. Household member i provides labor of type i . The wage rate and labor supply is set by a

union, described below. Households maximize

$$\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \left[\frac{(C_{h,t+j} - \gamma C_{h,t+j-1})^{1-\sigma}}{1-\sigma} - \frac{\chi}{1+\zeta} \int_0^1 N_{h,i,t+j}^{1+\zeta} di \right] \quad (47)$$

by choosing $C_{h,t}$, $b_{h,t}$, $u_{h,t}$, $I_{h,t}$, and $K_{h,t+1}$ subject to the budget constraints, for $j \geq 0$,

$$R_{t+j}^{-1} b_{h,t+j} + \frac{\vartheta}{2} b_{h,t+j}^2 + C_{h,t+j} + I_{h,t+j} + a(u_{h,t+j}) K_{h,t+j} \leq \int_0^1 (1 + \tau_w) w_{i,t+j} N_{h,i,t+j} di + r_{t+j}^K u_{h,t+j} K_{h,t+j} - \tau_{t+j} + \frac{b_{h,t+j-1}}{\pi_{t+j}} \quad (48)$$

and the capital accumulation equations

$$K_{h,t+j+1} = (1 - \delta) K_{h,t+j} + \left(1 - \frac{\phi_I}{2} \left(\frac{I_{h,t+j}}{I_{h,t+j-1}} - 1 \right)^2 \right) I_{h,t+j}. \quad (49)$$

If $\sigma = 1$, it is understood that the utility function for consumption takes the log form. Here, $C_{h,t}$ is household-specific consumption, $b_{h,t}$ are household-specific end-of-period real bond holdings, $u_{h,t}$ is the household-specific capital utilization rate, $I_{h,t}$ is household-specific investment, and $K_{h,t}$ is the beginning-of-period household-specific capital stock. R_t is the nominal interest rate on one-period nominal bonds, $w_{i,t}$ is the wage rate paid to labor-type i , r_t^K is the rental rate of capital, τ_t are taxes net profits, and $\pi_t \equiv P_t/P_{t-1}$. Note that capital and utilization provide capital services, $u_{h,t} K_{h,t}$. The term $a(u_{h,t}) K_{h,t}$ represents costs from utilization deviating from 1. In steady state, the form of $a(\cdot)$ ensures $u_{h,t} = 1$. The first-order conditions of household h are

$$\Lambda_{h,t} - (C_{h,t} - \gamma C_{h,t-1})^{-\sigma} + \gamma \mathbb{E}_t \beta (C_{h,t+1} - \gamma C_{h,t})^{-\sigma} = 0 \quad (50)$$

$$\Lambda_{h,t} + \vartheta R_t b_{h,t} \Lambda_{h,t} - R_t \beta \mathbb{E}_{h,t} \frac{1}{\pi_{t+1}} \Lambda_{h,t+1} = 0 \quad (51)$$

$$a'(u_{h,t}) - r_t^K = 0 \quad (52)$$

$$\begin{aligned} \Lambda_{h,t} - \left(1 - \frac{\phi_I}{2} \left(\frac{I_{h,t}}{I_{h,t-1}} - 1 \right)^2 \right) \eta_{h,t} + \phi_I \left(\frac{I_{h,t}}{I_{h,t-1}} - 1 \right) \frac{I_{h,t}}{I_{h,t-1}} \eta_{h,t} \\ - \beta \mathbb{E}_{h,t} \phi_I \left(\frac{I_{h,t+1}}{I_{h,t}} - 1 \right) \left(\frac{I_{h,t+1}}{I_{h,t}} \right)^2 \eta_{h,t+1} = 0 \end{aligned} \quad (53)$$

$$\eta_{h,t} + \beta \mathbb{E}_t a(u_{h,t+1}) \Lambda_{h,t+1} - \beta \mathbb{E}_t r_{t+1}^K u_{h,t+1} \Lambda_{h,t+1} - \beta \mathbb{E}_t \eta_{h,t+1} (1 - \delta) = 0 \quad (54)$$

Here, $\Lambda_{h,t}$ is the Lagrange multiplier on the budget constraint and $\eta_{h,t}$ is the Lagrange multiplier on the capital accumulation equation.

D.2 Monopolists

Monopolist f produces a differentiated good, $Y_{f,t}$. Goods are aggregated by competitive firms using the technology $Y_t = \left(\int_0^1 Y_{f,t}^{\frac{\varepsilon-1}{\varepsilon}} df \right)^{\frac{\varepsilon}{\varepsilon-1}}$. The price of Y_t is $P_t = \left(\int_0^1 P_{f,t}^{1-\varepsilon} df \right)^{\frac{1}{1-\varepsilon}}$. Monopolists face demand curves of the form given in equation (2). Production is done so that

$$Y_{f,t} = N_{f,t}^v \left(\tilde{K}_{f,t} \right)^{1-v} \quad (55)$$

where $\tilde{K}_{f,t}$ are capital services purchased by firm f . Notice that monopolists will solve a cost minimization problem that implies

$$\frac{w_t}{r_t^K} = \frac{v \tilde{K}_{f,t}}{(1-v) N_{f,t}} \quad (56)$$

$$s_t = \frac{w_t^v (r_t^K)^{1-v}}{v^v (1-v)^{1-v}}. \quad (57)$$

Monopolists update their price with probability $1 - \xi_P$. When they update their price, they choose $\tilde{P}_{f,t}$ to maximize the present-discount value of profits

$$\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \xi_P^j \Lambda_{t+j} \left(\frac{\tilde{P}_{f,t}}{P_{t+j}} - (1 - \tau_f) s_{t+j} \right) \left(\frac{\tilde{P}_t}{P_{t+j}} \right)^{-\varepsilon} Y_{t+j}. \quad (58)$$

Here

$$\Lambda_t - (C_t - \gamma C_{t-1})^{-\sigma} + \gamma \mathbb{E}_t \beta (C_{t+1} - \gamma C_t)^{-\sigma} = 0 \quad (59)$$

is the firm's perception of how they should discount flows to the household, s_{t+j} is real marginal cost in period $t+j$, and profits in period $t+j$ are discounted by an aggregate version of the household marginal utility of consumption. Assuming $(1 - \tau_f) \frac{\varepsilon}{\varepsilon-1} = 1$, the first-order condition is

$$\sum_{j=0}^{\infty} \beta^j \xi_P^j \mathbb{E}_t \Lambda_{t+j} \left(s_{t+j} - \tilde{p}_{f,t} \frac{P_t}{P_{t+j}} \right) \left(\frac{P_t}{P_{t+j}} \right)^{-\varepsilon} Y_{t+j} = 0 \quad (60)$$

where $\tilde{p}_{f,t} \equiv \tilde{P}_{f,t}/P_t$. Let

$$F_{P,t} = \sum_{j=0}^{\infty} \beta^j \xi_P^j \mathbb{E}_t \Lambda_{t+j} \left(\frac{P_t}{P_{t+j}} \right)^{1-\varepsilon} Y_{t+j} = \Lambda_t Y_t + \beta \xi_P \mathbb{E}_t \pi_{t+1}^{\varepsilon-1} F_{P,t+1} \quad (61)$$

Notice that we have assumed that the expectations of the firms satisfy the law of iterated expectations. Similarly, let

$$Z_{P,t} = \sum_{j=0}^{\infty} \beta^j \xi_P^j \mathbb{E}_t \Lambda_{t+j} \left(\frac{P_t}{P_{t+j}} \right)^{-\varepsilon} Y_{t+j} s_{t+j} = \Lambda_t Y_t s_t + \beta \xi_P \mathbb{E}_t \pi_{t+1}^\varepsilon Z_{P,t+1} \quad (62)$$

Then the first-order condition can be written as

$$\tilde{p}_{f,t} = \frac{Z_{P,t}}{F_{P,t}}. \quad (63)$$

D.3 Unions

We assume that each household has a continuum of members indexed by $i \in (0, 1)$. Each member consumes the same amount in every period (perfect consumption insurance). The index i indicates the labor type provided by the household member. Members of type i from all of the households belong to a union that sets the wage for type- i labor. With probability $1 - \xi_W$ a union is able to update the wage for its type. Each union provides all labor demanded at the posted wage and demands equal amounts of labor from each union member.

Labor is combined into a productive labor unit using the technology

$$N_t = \left(\int_0^1 \left(\int_0^1 N_{h,i,t} dh \right)^{\frac{\varepsilon-1}{\varepsilon}} di \right)^{\frac{\varepsilon}{\varepsilon-1}}$$

. A unit of N_t costs $W_t = \left(\int_0^1 W_{i,t}^{1-\varepsilon} di \right)^{\frac{1}{1-\varepsilon}}$. Demand for $N_{h,i,t}$ is given by equation (6). When a union sets its wage it does so by choosing $\tilde{W}_{i,t}$ to maximize equation (7).

We assume unions have the same expectation operator as households and firms. Additionally, unions take the demand curve for their labor type as given. Assuming $\frac{\varepsilon-1}{\varepsilon} (1 + \tau_w) = 1$, and defining $\tilde{w}_{i,t} \equiv \tilde{W}_{i,t}/P_t$ and $w_t \equiv W_t/P_t$, union i 's first-order condition is

$$\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \xi_W^j \left[\frac{\Lambda_{t+j} P_t w_t}{P_{t+j}} \left(\frac{\tilde{w}_{i,t}}{w_t} \right)^{1+\varepsilon\zeta} - \chi \left(\frac{w_t}{w_{t+j}} \frac{P_t}{P_{t+j}} \right)^{-\varepsilon\zeta} N_{t+j}^\zeta \right] \left(\frac{w_t}{w_{t+j}} \frac{P_t}{P_{t+j}} \right)^{-\varepsilon} N_{t+j} = 0 \quad (64)$$

Define

$$F_{W,t} = \Lambda_t N_t w_t + \beta \xi_W \mathbb{E}_t \pi_{t+1}^{\varepsilon-1} \left(\frac{w_t}{w_{t+1}} \right)^{1-\varepsilon} F_{W,t+1} \quad (65)$$

$$Z_{W,t} = \chi N_t^{1+\zeta} + \beta \xi_W \mathbb{E}_t \pi_{t+1}^{\varepsilon(1+\zeta)} \left(\frac{w_t}{w_{t+1}} \right)^{-\varepsilon(1+\zeta)} Z_{W,t+1}. \quad (66)$$

Because the law of iterated expectations holds for the subjective beliefs of unions, the first-order condition can be written as

$$F_{W,t} \left(\frac{\tilde{w}_{i,t}}{w_t} \right)^{1+\epsilon\zeta} = Z_{W,t}. \quad (67)$$

D.4 Aggregation and Government Policy

Aggregate output is given by

$$Y_t = C_t + I_t + G_t + \int_0^1 a(u_{h,t}) K_{h,t} dh. \quad (68)$$

Integrating over demand curves for intermediate goods and firm-specific production functions yields

$$\int_0^1 N_{f,t}^v \left(\tilde{K}_{f,t} \right)^{1-v} df = \int_0^1 \left(\frac{P_{f,t}}{P_t} \right)^{-\varepsilon} df Y_t. \quad (69)$$

Aggregate capital services are given by

$$\tilde{K}_t = \int_0^1 u_{h,t} K_{h,t} dh.$$

All firms pick the same capital-to-labor ratio, so

$$p_t^* N_t^v \left(\tilde{K}_t \right)^{1-v} = Y_t. \quad (70)$$

From the definitions of p_t^* and P_t ,

$$(p_t^*)^{-1} = (1 - \xi_P) \left(\frac{1 - \xi_P \pi_t^{\varepsilon-1}}{1 - \xi_P} \right)^{-\frac{\varepsilon}{1-\varepsilon}} + \xi_P \pi_t^{\varepsilon} (p_{t-1}^*)^{-1}. \quad (71)$$

Aggregate profits are given by

$$\int_0^1 \left(\frac{P_{f,t}}{P_t} - (1 - \tau_f) s_t \right) \left(\frac{P_{f,t}}{P_t} \right)^{-\varepsilon} Y_t df = \left(1 - (1 - \tau_f) \frac{s_t}{p_t^*} \right) Y_t. \quad (72)$$

We assume the government runs a balanced budget period by period. So,

$$\tau_t = G_t + \tau_w w_t N_t - \left(1 - \frac{s_t}{p_t^*} \right) Y_t. \quad (73)$$

We assume

$$G_t = G \quad (74)$$

and

$$R_t = (1 - \rho_R) \left(\beta^{-1} + \alpha_\pi (\pi_t - 1) + \alpha_Y \frac{Y_t - Y}{Y} \right) + \rho_R R_{t-1}. \quad (75)$$

D.5 REE Steady State

It is either the case that $R = \beta^{-1}$ and $\pi = 1$ or it is the case that $R = 1$ and $\pi = \beta$. We choose χ so that $N = 1$ when $\pi = 1$. We also set $G = 0.2Y$ where Y is the steady state value when $\pi = 1$. The form of the function $a(\cdot)$ ensures that $u = 1$ in steady state. See CEE (2005) for further discussion about $a(\cdot)$.

D.6 Beliefs and the Period-Equilibrium

Define

$$\begin{aligned} x_t &= [C_t, \pi_t, w_t, I_t]' , a_t = [C_t, \pi_t, w_t, I_t, p_t^*, R_t, Y_t, u_t, s_t, r_t^K, N_t, G_t, K_{t+1}, \tau_t]' \\ h_t &= [C_{h,t}, b_{h,t}, u_{h,t}, I_{h,t}, K_{h,t+1}, \Lambda_{h,t}, \eta_{h,t}, \tilde{p}_{f,t}, \pi_{f,t}, \tilde{w}_{i,t}, w_{i,t}, F_{P,t}, Z_{P,t}, F_{W,t}, Z_{W,t}, \Lambda_t]' \end{aligned}$$

We include Λ_t in h_t because it represents firms' and unions' perceptions of how to discount flow to the household. We assume that agents form beliefs as in equation (35). So that the bond market clears, we impose that it must be that $\hat{C}_{h,t} = \hat{C}_t$. So that the goods market clears, we impose that firm pricing decisions are consistent with aggregate inflation. To do that, we find it useful to define $\pi_{f,t}$ and $p_{f,t}^*$ so that

$$1 = (1 - \xi_P) \tilde{p}_{f,t}^{1-\epsilon} + \xi_P \pi_{f,t}^{\epsilon-1}. \quad (76)$$

We also need to impose that union wage-setting decisions are consistent with the aggregate real wage. To do that, we find it useful to define

$$w_{i,t}^{1-\epsilon} = (1 - \xi_W) \tilde{w}_{i,t}^{1-\epsilon} + \xi_W \pi_t^{\epsilon-1} w_{t-1}^{1-\epsilon}, \quad (77)$$

which follows from the definition of W_t . Finally, we need to impose that $I_{h,t} = I_t$ so that investment is consistent with overall aggregate output.

We write the restrictions we impose on the period equilibrium as $C_t = C_{h,t}$, $\pi_t = \pi_{f,t}$, $w_t = w_{i,t}$, $I_t = I_{h,t}$. It is now straightforward to put the model in the form discussed in Appendix B. Agents believe with certainty that the static equations (68), (70), (71), (73), (74), and (75), hold in each period. They also believe that the following equations hold with certainty $a'(u_t) = r_t^K$, $\frac{w_t}{r_t^K} = \frac{vu_t K_t}{(1-v)N_t}$, $s_t = \frac{w_t^v (r_t^K)^{1-v}}{v^v (1-v)^{1-v}}$, and $K_{t+1} = (1 - \delta) K_t + \left(1 - \frac{\phi_I}{2} \left(\frac{I_t}{I_{t-1}} - 1\right)^2\right) I_t$. Here, $K_t = \int_0^1 K_{h,t} dh$. These equations are aggregate versions of equations that hold in every period. These equation, along with the identities from x_t to a_t , define the mapping $a_t = F(x_t, a_{t-1})$ and their (log-)linear form can be expressed as in equation (36).

To construct equation (37), we use (log-)linear versions of equations (48), (49), (50), (51), (52), (53), (54), (59), (61), (62), (63), (65), (66), (67), (76), and (77).

D.7 Parameter Values

We use the following parameter values for our version of the CEE (2005) model with Calvo-style nominal rigidities.

Table 4: Parameter values for the CEE (2005) model

Parameter	Value	Description
β	$1.03^{-0.25}$	Rate of time discounting
γ	0.65	Habit formation parameter
ε	6	Governs demand elasticity in demand for goods
ϵ	21	Governs demand elasticity in demand for labor
ξ_P	0.6	Calvo parameter for prices
ξ_W	0.64	Calvo parameter for wages
σ	1	Log utility
ζ	1	Governs elasticity of labor supply
ρ_R	0.8	Interest rate smoothing
α_π	1.5	Coefficient on inflation in monetary policy rule
α_Y	0.25	Coefficient on output in monetary policy rule
ν	0.64	Cobb-Douglas parameter
$\frac{a''(1)}{a'(1)}$	0.01	Capital utilization parameter
δ	0.025	Capital depreciation parameter
ϕ_I	2.48	Investment adjustment cost parameter
$\frac{G}{Y}$	0.2	Government purchases in steady state
ϑ	0.000001	Bond-holding cost parameter

E Linearization When the ZLB is Binding

Here, we describe how we construct the period equilibrium in the simple NK model with only sticky prices and Rotemberg-style adjustment costs when the ZLB is binding. 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.

To fix notation, we let $\tilde{X}_t$ be value of the variable X_t when $r_t = r^\ell$ and let $\tilde{X}$ be its value in an REE. Let $\hat{\tilde{X}}_t = \log(\tilde{X}_t/\tilde{X})$. Note that we continue to use $\hat{b}_{h,t}$ and $\hat{p}_{f,t}$ 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 about $\tilde{C}_t$ and $\tilde{\pi}_t$ using

$$\hat{\tilde{C}}_t = \tilde{m}_{C,t-1} + \tilde{\sigma}_{\nu,C,t-1} \tilde{\nu}_{C,t} \quad (78)$$

$$\hat{\tilde{\pi}}_t = \tilde{m}_{\pi,t-1} + \tilde{\sigma}_{\nu,\pi,t-1} \tilde{\nu}_{\pi,t}. \quad (79)$$

Households and firms know that equations the following equations hold in every period: $Y_t = N_t$, $Y_t = (C_t + G_t) (1 + \frac{\Phi}{2}(\pi_t - 1)^2)$, $w_t = \chi N_t C_t$, $\tau_t = G_t - (1 - w_t)Y_t + \frac{\Phi}{2}(\pi_t - 1)^2(C_t + G_t)$, $G_t = G$, $R_t = \max \{1, \beta^{-1} + \alpha_\pi (\pi_t - 1) + \alpha_Y \frac{Y_t - Y}{Y}\}$. The Euler equation is

$$\frac{1}{\tilde{C}_{h,t} \tilde{R}_t} + \vartheta \tilde{b}_{h,t} = \frac{1}{1 + r_\ell} \mathbb{E}_t \frac{p}{\tilde{\pi}_{t+1} \tilde{C}_{h,t+1}} + \frac{1}{1 + r_\ell} \mathbb{E}_t \frac{1 - p}{\pi C_{h,t+1} (\tilde{b}_{h,t})}.$$

We emphasize that in the case that $r_{t+1} = r^{ss}$, we will have $\pi_{t+1} = \pi^{ss}$ and $C_{h,t+1}$ is a known function of household bond holdings. The firm's pricing optimality condition becomes

$$\begin{aligned} \left(\frac{\tilde{p}_{f,t}}{\tilde{p}_{f,t-1}} \tilde{\pi}_t - 1 \right) \tilde{\pi}_t \frac{\tilde{p}_{f,t}}{\tilde{p}_{f,t-1}} &= \frac{\varepsilon - 1}{\Phi_P} [\tilde{w}_t - \tilde{p}_{f,t}] \tilde{p}_{f,t}^{-\varepsilon} \frac{\tilde{Y}_t}{\tilde{C}_t + \tilde{G}_t} \\ &+ \frac{p}{1 + r_\ell} \mathbb{E}_t \left(\frac{\tilde{p}_{f,t+1}}{\tilde{p}_{f,t}} \tilde{\pi}_{t+1} - 1 \right) \frac{\tilde{p}_{f,t+1}}{\tilde{p}_{f,t}} \tilde{\pi}_{t+1} \left(\frac{\tilde{C}_{t+1} + \tilde{G}_{t+1}}{\tilde{C}_t + \tilde{G}_t} \right) \frac{\tilde{C}_t}{\tilde{C}_{t+1}} \\ &+ \frac{1 - p}{1 + r_\ell} \mathbb{E}_t \left(\frac{p_{f,t+1}}{\tilde{p}_{f,t}} \pi - 1 \right) \frac{p_{f,t+1}}{\tilde{p}_{f,t}} \pi \left(\frac{C + G}{\tilde{C}_t + \tilde{G}_t} \right) \frac{\tilde{C}_t}{C}. \end{aligned}$$

We can now proceed as in Appendix B to solve for the log-linear approximation to the period equilibrium. Note that we log-linearize $\tilde{X}_t$ around $\tilde{X}$. We need an approximation in which households and firms contemplate variables at the ZLB that grow slower than $\tilde{R}^{-1} \tilde{\pi} \tilde{\beta}^{-1} p^{-1}$. Because of the probability of switching from the ZLB to the steady state, the growth condition 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. Note that in any learning equilibrium, $\tilde{b}_{h,t} = b_{h,t} = 0$ and $\tilde{p}_{f,t} = p_{f,t} = 1$ in every period.

F Solution Algorithm for Nonlinear NK Model When the ZLB is Binding

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 $r_t = r^\ell$. Here, we maintain the assumption that households and firms know the REE in the steady state. With this solution in hand, we then turn to the period when $r_t = r^\ell$, which is where we consider learning. We find it useful to represent the model recursively in this appendix. This appendix explains how to solve the model under internalized learning. At the end of the appendix, we explain how to solve the model using that solution to generate

results for anticipated utility.

F.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.

F.1.1 Household Problem

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

$$V_h^{ss}(b_{h,t-1}) = \max_{C_{h,t}, N_{h,t}, b_{h,t}} \left\{ \log(C_{h,t}) - \frac{\chi}{2} (N_{h,t})^2 + \beta V_h^{ss}(b_{h,t}) \right\}$$

subject to

$$C_{h,t} + \frac{b_{h,t}}{R^{ss}} \leq \frac{b_{h,t-1}}{\pi^{ss}} + w^{ss} N_{h,t} - \tau^{ss}.$$

Here, $b_{h,t-1}$ and $b_{h,t}$ are household h 's real bond holdings chosen in the previous and current period, respectively. The variables $C_{h,t}$ and $N_{h,t}$ are household h 's consumption and labor supply. The aggregate variable R^{ss} , π^{ss} , w^{ss} , and τ^{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 $b_{h,t} \in [\underline{b}, \bar{b}]$. Implicitly, we have functions $C_{h,t} = C_h^{ss}(b_{h,t-1})$, $N_{h,t} = N_h^{ss}(b_{h,t-1})$, and $b_{h,t} = b_h^{ss}(b_{h,t-1})$. Assuming the constraint on $b_{h,t}$ is not binding, household maximization implies

$$\frac{1}{C_h^{ss}(b_{h,t-1})} = \beta R^{ss} \frac{1}{C_h^{ss}(b_{h,t}) \pi^{ss}} \quad (80)$$

$$\chi C_h^{ss}(b_{h,t-1}) N_h^{ss}(b_{h,t-1}) = w^{ss} \quad (81)$$

$$C_h^{ss}(b_{h,t-1}) + \frac{b_h^{ss}(b_{h,t-1})}{R^{ss}} = \frac{b_{h,t-1}}{\pi^{ss}} + w^{ss} N_h^{ss}(b_{h,t-1}) - \tau^{ss} \quad (82)$$

We define a grid over $[\underline{b}, \bar{b}]$ and approximate the functions $C_h^{ss}(b_{h,t-1})$, $N_h^{ss}(b_{h,t-1})$, and $b_h^{ss}(b_{h,t-1})$ on that grid in the following way.⁴²

- (i) We conjecture a value for $b_h^{ss}(b_{h,t-1})$ at each grid point, which we denote by $b_h^{ss,i}(b_{h,t-1})$.
- (ii) Note that equations (81) and (82) can be written as

$$\chi C_h^{ss}(b_{h,t-1}) \left(C_h^{ss}(b_{h,t-1}) + \frac{b_h^{ss,i}(b_{h,t-1})}{R^{ss}} - \frac{b_{h,t-1}}{\pi^{ss}} + \tau^{ss} \right) = (w^{ss})^2.$$

⁴²In our implementation, we set $-\underline{b} = \bar{b} = 1$, 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 $b_{h,t-1} = b_{h,t} = 0$ in both REE and in learning equilibria.

The left-hand-side is increasing in $C_h^{ss}(b_{h,t-1}) \geq 0$. For every $b_{h,t-1}$, we solve for the value of $C_h^{ss}(b_{h,t-1})$ that makes this hold with equality. We call this the conjectured value for $C_h^{ss}(b_{h,t-1})$ and denote it by $C_h^{ss,i}(b_{h,t-1})$. Note that with $C_h^{ss,i}(b_{h,t-1})$, we can back out $N_h^{ss,i}(b_{h,t-1})$ from equation (81).

- (iii) For each grid point, $b_{h,t-1}$, find $b_{h,t}$ that solves the following version of equation (80)

$$C_{h,t}\beta R^{ss} \frac{1}{C_h^{ss,i}(b_{h,t})} - 1 = 0$$

where $C_{h,t} \geq 0$ solves

$$\chi C_{h,t} \left(C_{h,t} + \frac{b_{h,t}}{R^{ss}} - \frac{b_{h,t-1}}{\pi^{ss}} + \tau^{ss} \right) = (w^{ss})^2.$$

We use linear interpolation to compute $C_h^{ss,i}(b_{h,t})$ for values of $b_{h,t}$ that fall between grid points. If the procedure would set $b'_h > \bar{b}$ or $b'_h < \underline{b}$, we set $b_{h,t}$ to the respective endpoint of the grid. We record the value of $b_{h,t}$ by updating the conjectured function $b_h^{ss}(\cdot)$ using $b_h^{ss,i+1}(b_{h,t-1}) = b_{h,t}$.

- (iv) Having computed $b_h^{ss,i+1}(b_{h,t-1})$ for every grid point, we check to see if

$$|b_h^{ss,i+1}(b_{h,t-1}) - b_h^{ss,i}(b_{h,t-1})| < \epsilon$$

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 $b_h^{ss,i}(b_{h,t-1}) = b_h^{ss,i+1}(b_{h,t-1})$ at each grid point and repeat steps (ii), (iii), and (iv).

Because $\beta \frac{R^{ss}}{\pi^{ss}} = 1$, it is not surprising that we find that $b_h^{ss}(b_{h,t-1}) = b_{h,t-1}$.

F.1.2 Firm Problem

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

$$V_f^{ss}(p_{f,t-1}) = \max_{p_{f,t}} \left\{ \frac{1}{C^{ss}} (p_{f,t} - (1 - \tau_f) w^{ss}) (p_{f,t})^{-\varepsilon} Y^{ss} - \frac{1}{C^{ss}} \frac{\Phi_P}{2} \left(\frac{p_{f,t}}{p_{f,t-1}} \pi^{ss} - 1 \right)^2 (C^{ss} + G^{ss}) + \beta V_f^{ss}(p_{f,t}) \right\}.$$

Here, $p_{f,t-1}$ and $p_{f,t}$ 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 $\log(p_{f,t}) \in [\underline{p}, \bar{p}]$.⁴³ Implicitly, we have a function

⁴³In our implementation, we set $-\underline{p} = \bar{p} = 1$. We use a symmetric grid with 25 points that includes zero that places more points near zero than at more extreme values because $\log(p_{f,t-1}) = \log(p_{f,t}) = 0$ in both REE and in learning equilibria.

$p_f^{ss}(p_{f,t-1})$. Assuming the constraint on $\log(p_f)$ is not binding, firm optimality implies

$$\begin{aligned} \Phi_P \left(\frac{p_f^{ss}(p_{f,t-1})}{p_{f,t-1}} \pi^{ss} - 1 \right) \frac{1}{p_f} \pi^{ss} (C^{ss} + G^{ss}) &= (\varepsilon - 1) \left(\frac{w^{ss}}{p_f^{ss}(p_{f,t-1})} - 1 \right) (p_f^{ss}(p_{f,t-1}))^{-\varepsilon} Y^{ss} \\ + \beta \Phi_P \left(\frac{p_f^{ss}(p_{f,t-1})}{p_f^{ss}(p_{f,t-1})} \pi^{ss} - 1 \right) \frac{p_f^{ss}(p_{f,t-1})}{(p_f^{ss}(p_{f,t-1}))^2} \pi^{ss} (C^{ss} + G^{ss}) \end{aligned} \quad (83)$$

We define a grid over $[\underline{p}, \bar{p}]$ and approximate the function $p_f^{ss}(p_{f,t-1})$ on that grid in the following way.

- (i) We conjecture a value for $p_f^{ss}(p_{f,t-1})$ at each grid point. Call the conjectured value $p_f^{ss,i}(p_{f,t-1})$.
- (ii) For each grid point, $p_{f,t-1}$, find $p_{f,t}$ that solves the following version of equation (83)

$$\begin{aligned} \Phi_P \left(\frac{p_{f,t}}{p_{f,t-1}} \pi^{ss} - 1 \right) \frac{1}{p_{f,t-1}} \pi^{ss} (C^{ss} + G^{ss}) &= (\varepsilon - 1) \left(\frac{w^{ss}}{p_{f,t}} - 1 \right) (p_{f,t})^{-\varepsilon} Y^{ss} \\ + \beta \Phi_P \left(\frac{p_f^{ss,i}(p_{f,t})}{p_{f,t}} \pi^{ss} - 1 \right) \frac{p_f^{ss,i}(p_{f,t})}{(p_{f,t})^2} \pi^{ss} (C^{ss} + G^{ss}) \end{aligned}$$

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

- (iii) Having computed $p_f^{ss,i+1}(p_{f,t-1})$ for every grid point, we check to see if

$$|p_f^{ss,i+1}(p_{f,t-1}) - p_f^{ss,i}(p_{f,t-1})| < \epsilon$$

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 $p_f^{ss,i}(p_{f,t-1}) = p_f^{ss,i+1}(p_{f,t-1})$ at each grid point and repeat steps (ii) and (iii).

F.2 Solution when $r_t = r^\ell$

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

F.2.1 Beliefs

Before presenting the household and firm problems, some comments about beliefs are in order when $r_t = r^\ell$. 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 $r_t = r^\ell$, $\log(C_t)$ and $\log(\pi_t)$, have uncorrelated Normal distributions with unknown means and variances. That is, $\log(\pi_t) \sim N(\mu_\pi, \sigma_\pi^2)$ and $\log(C_t) \sim N(\mu_C, \sigma_C^2)$. We assume that households and firms have beliefs about the means and variances of the distributions for $\log(C_t)$ and $\log(\pi_t)$ that are characterized by density functions that are proportional to Normal-inverse-gamma distributions. (See Zellner (1971) and Murphy (2007) for a discussion of the Normal-inverse-gamma distribution). 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, $\mathbb{E}[\pi_t] = \mathbb{E}[C_t] = \infty$, which would challenge the applicability of an expected utility framework (see Geweke (2001) for related discussion). The distributions characterizing beliefs are independent across C_t and π_t . That is, for $i \in \{\pi, C\}$, $\mu_i \in (-\infty, \infty)$ and $\sigma_i^2 \in [0, \bar{\sigma}_i^2]$ we have

$$\Pr(\sigma_i^2 | \alpha_{i,t-1}, \beta_{i,t-1}) = \frac{\frac{\beta_i^{\alpha_{i,t-1}}}{\Gamma(\alpha_{i,t-1})} \left(\frac{1}{\sigma_i^2}\right)^{\alpha_{i,t-1}+1} \exp\left(-\frac{\beta_{i,t-1}}{\sigma_i^2}\right)}{\frac{\Gamma(\alpha_{i,t-1}, \beta_{i,t-1}/\bar{\sigma}_i^2)}{\Gamma(\alpha_{i,t})}},$$

$$\Pr(\mu_i | \sigma_i^2, m_{i,t-1}, \lambda_{i,t-1}) = \frac{\sqrt{\lambda_{i,t-1}}}{\sqrt{2\pi\sigma_i^2}} \exp\left(-\frac{\lambda_{i,t-1}}{2\sigma_i^2} (\mu_i - m_{i,t-1})^2\right).$$

See 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 prior of σ_i^2 is that $\mathbb{E}[\pi_t] < \infty$ if $\bar{\sigma}_\pi^2 < \infty$ and $\mathbb{E}[C_t] < \infty$ if $\bar{\sigma}_C^2 < \infty$.

Even though we truncate the distributions for σ_i^2 , we maintain the usual recursive updating equations because the likelihoods associated with observations of π_t and C_t are not truncated. Beliefs about consumption and inflation are parameterized by four values, $\alpha_{i,t-1}$, $\beta_{i,t-1}$, $m_{i,t-1}$, and $\lambda_{i,t-1}$ for $i \in \{\pi, C\}$. So, we have 8 total parameters. The recursive updating formulas (see Murphy (2007)) for these variables are $\lambda_{i,t} = \lambda_{i,t-1} + 1$, $\alpha_{i,t} = \alpha_{i,t-1} + 1/2$,

$$m_{i,t} = \frac{\lambda_{i,t-1}m_{i,t-1} + \log(i_t)}{\lambda_{i,t-1} + 1}, \quad \beta_{i,t} = \beta_{i,t-1} + \frac{\lambda_{i,t-1}(\log(i_t) - m_{i,t-1})^2}{2(\lambda_{i,t-1} + 1)}.$$

We need to include variables in Θ_{t-1} that will fully capture the values $\alpha_{i,t-1}$, $\beta_{i,t-1}$, $m_{i,t-1}$, and $\lambda_{i,t-1}$ for $i \in \{\pi, C\}$. First, we keep $\frac{1}{t_\ell}$ in Θ_{t-1} , which is the inverse of the 1 plus number of periods that r_t has been equal to r_ℓ up to period $t-1$. We keep the inverse because it is bounded between zero and one. From this value, we can trivially back out $\lambda_{i,t-1}$ and $\alpha_{i,t-1}$, given their values in the first period when $r_t = r_\ell$. We set the initial value of $\lambda_{i,t-1} = 1$ and the initial value of $\alpha_{i,t-1} = 2$. We keep $m_{C,t-1}$ and $m_{\pi,t-1}$ in Θ_{t-1} . And we also keep

$$\psi_{i,t} = \sqrt{\psi_{i,t-1}^2 \frac{2\alpha_{i,t}}{2\alpha_{i,t} + 1} + \frac{\lambda_{i,t-1}}{(\lambda_{i,t-1} + 1)} \frac{1}{2\alpha_{i,t} + 1} (\log(i_t) - m_{i,t-1})^2}.$$

Note that by setting $\beta_{i,t} = (\psi_{i,t})^2 \alpha_{i,t+1}$ it is clear that we recover the exactly recursive structure of β_i (given above). An advantage of using $\psi_{i,t}$ in Θ_t rather than $\beta_{i,t}$ is that $\psi_{i,t}$ is a consistent estimator for the standard deviation, whereas $\beta_{i,t}$ generally grows without bound (except when the standard deviation is zero). Keeping the values of Θ_t within bounded grids will be important for the purposes of approximation. In total, $\Theta_t = [\frac{1}{t^\ell}, m_{\pi,t}, m_{C,t}, \psi_{\pi,t}, \psi_{C,t}]$ has five elements (where t^ℓ is the number of periods that the ZLB has persisted) and we have a mapping from Θ_{t-1} to $\alpha_{i,t}$, $\beta_{i,t}$, $m_{i,t}$, and $\lambda_{i,t}$ for $i \in \{\pi, C\}$. We also have a law of motion for Θ_t so that $\Theta_t = L(\Theta_{t-1}, [\pi_t, C_t]')$.

An advantage of the Normal-inverse-gamma setup detailed above is that we can have analytic expressions for the distribution for the variables $\log(\pi_t)$ and $\log(C_t)$ conditional on Θ_{t-1} . (See Zellner (1971) and Murphy (2007) for related discussion. See Bakshi and Skoulakis (2010) for a discussion of truncating the distribution of the variance.) In particular

$$\begin{aligned} \Pr(\log(i_t) | \Theta_{t-1}) &= \left(\frac{\lambda_{i,t-1} \alpha_{i,t-1}}{\beta_{i,t-1} (\lambda_{i,t-1} + 1)} \right)^{1/2} \frac{\Gamma(\alpha_{i,t-1} + 1/2)}{\Gamma(\alpha_{i,t-1}) \sqrt{2\pi\alpha_{i,t-1}}} \\ &\quad \times \left(1 + \frac{1}{2\alpha_{i,t-1}} \frac{(\log(i_t) - m_{i,t-1})^2}{\left(\frac{\lambda_{i,t-1} \alpha_{i,t-1}}{\beta_{i,t-1} (\lambda_{i,t-1} + 1)} \right)^{-1}} \right)^{-\alpha_{i,t-1} - 1/2} \frac{\kappa_{i,t}}{\kappa_{i,t-1}}, \end{aligned} \quad (84)$$

where $\kappa_{i,t} = \Gamma(\alpha_{i,t}, \beta_{i,t}/\bar{\sigma}_i^2) [\Gamma(\alpha_{i,t})]^{-1}$. Notice that $\kappa_{i,t}$ depends on the point of evaluation for $\log(i_t)$. If we ignored the ratio $\kappa_{i,t}/\kappa_{i,t-1}$, which would be correct in the case when $\bar{\sigma}_i^2 = \infty$, the pdf for $\log(i_t)$ is a t distribution with location parameter $m_{i,t-1}$, scale parameter $\left(\frac{\lambda_{i,t-1} \alpha_{i,t-1}}{\beta_{i,t-1} (\lambda_{i,t-1} + 1)} \right)^{-1/2}$, and $2\alpha_{i,t-1}$ degrees of freedom. If $\bar{\sigma}_i^2$ is large, $\kappa_{i,t}/\kappa_{i,t-1} \neq 1$ but is close to unity. For finite $\bar{\sigma}_i^2$, the ratio $\kappa_{i,t}/\kappa_{i,t-1}$ serves to thin the tails of the distribution of $\log(i_t)$ by down-weighting the probability of extreme values for $\log(i_t)$.⁴⁴ Because the density function of the t distribution is readily available and reliably computed in statistical software and because $\kappa_{i,t-1}$ and $\kappa_{i,t}$ are easily computed using readily available implementations of the gamma and incomplete gamma functions, we can use equation (84) for quadrature weighting. We use Gauss-Hermite quadrature with seven nodes to approximate integrals based on equation (84).

F.2.2 Household Problem

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

⁴⁴We set $\bar{\sigma}_i^2$ equal to the squared maximum value on the grid for $\psi_{i,t}$ (described below).

$$V_h(b_{h,t-1}, \Theta_{t-1}, x_t) = \max_{C_{h,t}, N_{h,t}, b_{h,t}} \left\{ \log(C_{h,t}) - \frac{\chi}{2} (N_{h,t})^2 \right. \\ \left. + \frac{1}{1+r^\ell} [p\mathbb{E}_t V_h(b_{h,t}, \Theta_t, x_{t+1}) + (1-p)V_h^{ss}(b_{h,t})] \right\}$$

subject to $C_{h,t} + \frac{b_{h,t}}{R_t} \leq \frac{b_{h,t-1}}{\pi_t} + w_t N_{h,t} - \tau_t$. Here, $x_t = [\pi_t, C_t]'$, $V_h^{ss}(\cdot)$ is the steady state value function for the household, which is defined above, and $\mathbb{E}_t$ denotes expectations of the household computed conditional on Θ_{t-1} and x_t . Given x_t and Θ_{t-1} , we have $\Theta_t = L(\Theta_{t-1}, x_t)$. So, the expectation of the household is taken with respect to x_{t+1} , 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 $Y_t = (C_t + G_t)(1 + \frac{\Phi_P}{2}(\pi_t - 1)^2)$, $N_t = Y_t$, and $w_t = \chi C_t Y_t$. Given x_t , with these assumptions R_t , w_t , and τ_t can be computed. The steady state values of aggregate variables are known to the household. We constrain households so that $b_{h,t} \in [\underline{b}, \bar{b}]$. The household optimization problem gives us implicit functions for $C_{h,t} = C_h(b_{h,t-1}, \Theta_{t-1}, x_t)$, $N_{h,t} = N_h(b_{h,t-1}, \Theta_{t-1}, x_t)$, and $b_{h,t} = b_h(b_{h,t-1}, \Theta_{t-1}, x_t)$. Considering interior solutions for $b_{h,t}$, we have

$$\frac{1}{C_h(b_{h,t-1}, \Theta_{t-1}, x_t)} = \frac{1}{1+r^\ell} R \left[p\mathbb{E}_t \left\{ \frac{1}{\pi_{t+1} C_h(b_{h,t}, \Theta_t, x_{t+1})} \right\} \right. \\ \left. + (1-p) \frac{1}{\pi^{ss} C_h^{ss}(b_{h,t})} \right] \quad (85)$$

$$w_t = \chi C_h(b_{h,t-1}, \Theta_{t-1}, x_t) N_h(b_{h,t-1}, \Theta_{t-1}, x_t) \quad (86)$$

$$C_h(b_{h,t-1}, \Theta_{t-1}, x_t) = \frac{b_{h,t-1}}{\pi_t} + w_t N_h(b_{h,t-1}, \Theta_{t-1}, x_t) - \tau_t - \frac{b_h(b_{h,t-1}, \Theta_{t-1}, x_t)}{R_t}. \quad (87)$$

Instead of approximating $C_h(b_{h,t-1}, \Theta_{t-1}, x_t)$, $N_h(b_{h,t-1}, \Theta_{t-1}, x_t)$, and $b_h(b_{h,t-1}, \Theta_{t-1}, x_t)$ directly, we approximate

$$v_h(b_{h,t}, \Theta_t) = \mathbb{E}_t \left\{ \frac{1}{\pi_{t+1} C_h(b_{h,t}, \Theta_t, x_t)} \right\}.$$

We take this approach because it eliminates x_t as a state variable in the approximation. We define grids on the elements of Θ_t and use the grid defined for $b_{h,t}$ in the steady state. We then approximate $v_h(b_{h,t}, \Theta_t)$ in the following way.⁴⁵

(i) We conjecture a value for $v_h(b_{h,t-1}, \Theta_{t-1})$ at each grid point in the cross product of the

⁴⁵The grids for $m_{i,t}$ contain 12 points that are not evenly spaced. They include each REE point as well as the target-inflation steady state. The remaining points are bunched relatively close to the REE points. The grid for $\psi_{C,t}$ contains 11 points that are evenly spaced from 0 to 0.1. The grid for $\psi_{\pi,t}$ 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.

- grids over the elements of $b_{h,t-1}$ and Θ_{t-1} .⁴⁶ Call the conjectured value $v_h^i(b_{h,t-1}, \Theta_{t-1})$.
- (ii) For a given grid point we use quadrature to get a value for $\mathbb{E}_{t-1} \{(\pi_t C_{h,t})^{-1}\}$. To solve for the expectation of interest, we need to solve for $C_{h,t}$ given many different values for x_t . Conditional on a value for x_t , equations (86) and (87) can be written as

$$\chi C_{h,t} \left(C_{h,t} + \frac{b_{h,t}}{R_t} - \frac{b_{h,t-1}}{\pi_t} + \tau_t \right) = w_t^2$$

The left-hand-side is increasing in $C_{h,t} \geq 0$. For a given $b_{h,t}$, we solve for the value of $C_{h,t}$ that makes this hold with equality. We then search for the value of $b_{h,t}$ that makes the following version of equation (85) hold with equality:

$$\frac{1}{C_{h,t}} = \frac{1}{1+r^\ell} R \left[p v_h^i(b_{h,t}, \Theta_t) + (1-p) \frac{1}{\pi^{ss} C_h^{ss}(b_{h,t+1})} \right].$$

We use linear interpolation to compute $v_h^i(b_{h,t}, \Theta_t)$ for values of $b_{h,t}$ and Θ_t that fall between grid points. If the procedure would set $b_{h,t} > \bar{b}$ or $b_{h,t} < \underline{b}$, we set $b_{h,t}$ to the respective endpoint of the grid for $b_{h,t}$. We record the associated value of $C_{h,t}$ and use it in the quadrature to compute $v_h^{i+1}(b_{h,t-1}, \Theta_{t-1}) = \mathbb{E}_{t-1} \{(\pi_t C_{h,t})^{-1}\}$.

- (iii) Having computed $v_h^{i+1}(b_{h,t-1}, \Theta_{t-1})$ for every grid point, we check to see if

$$|v_h^i(b_{h,t-1}, \Theta_{t-1}) - v_h^{i+1}(b_{h,t-1}, \Theta_{t-1})| < \epsilon$$

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

The grid that we use on $\frac{1}{t^\ell}$ is special. In particular, we let that grid be $[0, \frac{1}{99}, \frac{1}{98}, \dots, 1]$. 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 $\Theta_t = \Theta_{t-1}$ because $m_{i,t-1}$ and $\psi_{i,t-1}$ 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 $v_h^i(b_{h,t-1}, \Theta_{t-1})$ when $t^\ell = 99$ to evaluate $v_h^i(b_{h,t}, \Theta_t)$.⁴⁷ That is, when $t^\ell = 99$ we have to find a fixed point of this interpolation, which is computationally intense. When $(t^\ell)^{-1} = 98$, having approximated $v_h(b_{h,t}, \Theta_t)$ for $(t^\ell)^{-1} = 99$ means that can evaluate $v_h(b_{h,t}, \Theta_t)$ exactly at $(t^\ell)^{-1} = 99$ without reference to $v_h^i(b_{h,t}, \Theta_t)$. We approximate for $v_h(b_{h,t}, \Theta_t)$ when $(t^\ell)^{-1} = 98$ and work work back in

⁴⁶There are $435,600 = 12 \times 12 \times 11 \times 11 \times 25$ points in the cross product of the grids for $m_{i,t}$, $\psi_{i,t}$, and $b_{h,t}$. t_ℓ^{-1} is handled in a way discussed below.

⁴⁷When solving their model, Collin-Dufresne et al. (2016) assume that at some point in time the true parameters become known.

this way to $(t^\ell)^{-1} = 1$. This strategy fits this into the structure of steps 1-3 because we know that the value of $v_h(b_{h,t-1}, \Theta_{t-1})$ will not depend on its value at any t^ℓ that is smaller than implied by Θ_{t-1} . So, we have a block dependent structure to $v_h(b_{h,t-1}, \Theta_{t-1})$. Additionally, we know that t^ℓ will only take integer values.

F.2.3 Firm Problem

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

$$V_f(p_{f,t-1}, \Theta_{t-1}, x_t) = \max_{p'_f} \left\{ \frac{1}{C} \left((p_{f,t} - (1 - \tau_f) w_t) (p_{f,t})^{-\varepsilon} Y_t - \frac{\Phi_P}{2} \left(\frac{p_{f,t}}{p_{f,t-1}} \pi_t - 1 \right)^2 (C_t + G_t) \right) \right. \\ \left. + \frac{1}{1 + r^\ell} [p \mathbb{E}_t V_f(p_{f,t}, \Theta_t, x_{t+1}) + (1 - p) V_f^{ss}(p_{f,t})] \right\}$$

Here, $x_t = [\pi_t, C_t]'$, $V_f^{ss}(\cdot)$ is the steady state value function for the firm, which is defined above, and $\mathbb{E}_t$ denotes expectations of the firm conditional on Θ_{t-1} and x_t . Given x_t and Θ_{t-1} , we have $\Theta_t = L(\Theta_{t-1}, x_t)$. So, the expectation of the firm is taken with respect to x_{t+1} , 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 $Y_t = (C_t + G_t) (1 + \frac{\Phi_P}{2} (\pi_t - 1)^2)$, $N_t = Y_t$, and $w_t = \chi C_t Y_t$. Given x_t , with these assumptions w_t , G_t , and Y_t can be computed. The steady state values of aggregate variables are known to the firm. We constrain firms so that $\log(p_{f,t}) \in [\underline{p}, \bar{p}]$. Implicitly, from firm optimization we have a function $p_{f,t} = p_f(p_{f,t-1}, \Theta_{t-1}, x_t)$. Considering interior solutions for $p_{f,t}$, firm maximization implies

$$\Phi_P \left(\frac{p_f(p_{f,t-1}, \Theta_{t-1}, x_t)}{p_{f,t-1}} \pi_t - 1 \right) \frac{1}{p_{f,t-1}} \pi_t (C_t + G_t) = \\ (\varepsilon - 1) \left(\frac{w_t}{p_f(p_{f,t-1}, \Theta_{t-1}, x_t)} - 1 \right) (p_f(p_{f,t-1}, \Theta_{t-1}, x_t))^{-\varepsilon} Y_t \\ + \frac{1}{1 + r^\ell} p \mathbb{E}_t \frac{C_t}{C_{t+1}} \Phi_P \left(\frac{p_f(p_f(p_{f,t-1}, \Theta_{t-1}, x_t), \Theta_t, x_{t+1}))}{p_f(p_{f,t-1}, \Theta_{t-1}, x_t)} \pi_{t+1} - 1 \right) \frac{p_f(p_f(p_{f,t-1}, \Theta_{t-1}, x_t), \Theta_t, x_{t+1}))}{(p_f(p_{f,t-1}, \Theta_{t-1}, x_t))^2} \pi' (C' + G') \\ + \frac{1}{1 + r^\ell} \frac{C_t}{C^{ss}} (1 - p) \Phi_P \left(\frac{p_f^{ss}(p_f(p_{f,t-1}, \Theta_{t-1}, x_t))}{p_f(p_{f,t-1}, \Theta_{t-1}, x_t)} \pi^{ss} - 1 \right) \frac{p_f^{ss}(p_f(p_{f,t-1}, \Theta_{t-1}, x_t))}{(p_f(p_{f,t-1}, \Theta_{t-1}, x_t))^2} \pi^{ss} (C^{ss} + G^{ss}). \quad (88)$$

Instead of approximating $p_f(p_{f,t-1}, \Theta_{t-1}, x_t)$ directly, we approximate

$$v_f(p_{f,t}, \Theta_t) = \mathbb{E}_t \left\{ \frac{1}{C_{t+1}} \Phi_P \left(\frac{p_{f,t+1}}{p_{f,t}} \pi_{t+1} - 1 \right) \frac{p_{f,t+1}}{p_{f,t}} \pi_{t+1} (C_{t+1} + G_{t+1}) \right\}.$$

We take this approach because we can eliminate x_t as a state variable in the approximation. We use the same grids on the elements of Θ_{t-1} that we use for the household problem and the grid defined for $\log(p_{f,t-1})$ in the steady state and we approximate $v_f(p_{f,t-1}, \Theta_{t-1})$ in the following way.

- (i) We conjecture a value for $v_f(p_{f,t-1}, \Theta_{t-1})$ at each grid point in the cross product of the grids over the elements of $p_{f,t-1}$ and Θ_{t-1} . Call the conjectured value $v_f^i(p_{f,t-1}, \Theta_{t-1})$.
- (ii) For a given grid point we use quadrature to get a value for

$$\mathbb{E}_{t-1} \left\{ \frac{1}{C_t} \Phi_P \left(\frac{p_{f,t}}{p_{f,t-1}} \pi_t - 1 \right) \frac{p_{f,t}}{p_{f,t-1}} \pi_t (C_t + G_t) \right\}.$$

To solve for the expectation of interest, we need to solve for $p_{f,t}$ given many different values for x_t . Conditional on a value for x_t , we find a value of $p_{f,t}$ that solves the following version of equation (88)

$$\begin{aligned} \Phi_P \left(\frac{p_{f,t}}{p_{f,t-1}} \pi_t - 1 \right) \frac{1}{p_{f,t-1}} \pi_t (C_t + G_t) = \\ (\varepsilon - 1) \left(\frac{w_t}{p_{f,t}} - 1 \right) (p_{f,t})^{-\varepsilon} Y_t + \frac{1}{1 + r_\ell} p v_f^i(p_{f,t}, \Theta_t) \frac{C_t}{p_{f,t}} \\ + \frac{1}{1 + r_\ell} \frac{C_t}{C^{ss}} (1 - p) \Phi_P \left(\frac{p_f^{ss}(p_{f,t})}{p_{f,t}} \pi^{ss} - 1 \right) \frac{p_f^{ss}(p_{f,t})}{(p_{f,t})^2} \pi^{ss} (C^{ss} + G^{ss}). \end{aligned}$$

We use linear interpolation over $\log(p_{f,t})$ to compute $v_f^i(p_{f,t}, \Theta_t)$ for values of $\log(p_{f,t})$ and Θ_t that fall between grid points. If the procedure would set $\log(p_{f,t}) > \bar{p}$ or $\log(p_{f,t}) < \underline{p}$, we set $p_{f,t}$ to the respective endpoint of the grid for p_f . We record the value of $p_{f,t}$ and the associated aggregate variables so that the quadrature procedure can approximate

$$v_f^{i+1}(p_{f,t-1}, \Theta_{t-1}) = \mathbb{E}_{t-1} \left\{ \frac{1}{C_t} \Phi_P \left(\frac{p_{f,t}}{p_{f,t-1}} \pi_t - 1 \right) \frac{p_{f,t}}{p_{f,t-1}} \pi_t (C_t + G_t) \right\}.$$

- (iii) Having computed $v_f^{i+1}(p_{f,t-1}, \Theta_{t-1})$ for every grid point, we check to see if

$$|v_f^i(p_{f,t-1}, \Theta_{t-1}) - v_f^{i+1}(p_{f,t-1}, \Theta_{t-1})| < \epsilon$$

at every grid point for some small ϵ . If yes, we say that we have solved the household problem when $r_t = r_\ell$. If no, we set $v_f^i(p_{f,t-1}, \Theta_{t-1}) = v_f^{i+1}(p_{f,t-1}, \Theta_{t-1})$, 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 $(t^\ell)^{-1}$.

F.3 Learning Equilibria

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

- (i) Set $r_t = r_\ell$ and assume a value for Θ_{t-1} for $t = 1$.
- (ii) Conjecture a value for π_t .

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

$$\frac{1}{C_t} = \frac{1}{1+r^\ell} R_t \left[pv_h(0, L(\Theta_{t-1}, [\pi_t, C_t])) + (1-p) \frac{1}{\pi^{ss} C_h^{ss}(0)} \right].$$

(b) Check to see if the following equation holds

$$\begin{aligned} \Phi_P(\pi_t - 1) \pi_t (C_t + G_t) &= (\varepsilon - 1) (w_t - 1) + \frac{1}{1+r_\ell} pv_f(1, L(\Theta_{t-1}, [\pi_t, C_t])) C_t \\ &+ \frac{1}{1+r_\ell} \frac{C_t}{C^{ss}} (1-p) \Phi_P(\pi^{ss} - 1) \pi^{ss} (C^{ss} + G^{ss}). \end{aligned}$$

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 .

(iii) Set $\Theta_{t+1} = L(\Theta_t, [\pi_t, C_t])$ 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. Because it is much less computationally demanding, the anticipated utility solution can also be calculated directly using $m_{i,t-1}$ and $\psi_{i,t-1}$ without relying on interpolation between grid points of those variables.

F.4 Forward Guidance

When we implement forward guidance, we assume that households and firms know the aggregate values of consumption and inflation in period I (the period that r_t changes back to r^{ss} but R_t remains at the ZLB). We denote these quantities by C^I and π^I . We also assume that households and firms know how they will act in period I . However, they are unsure about the period when $r_t = r^\ell$. They have to learn about the data generating process for C_t and π_t while $r_t = r^\ell$. With these assumptions, the model can now be cast as in the previous three sub-sections, but with householding and firms thinking that with probability $1-p$ they will be in period I and using the value functions and decision rules in period I .

G REE Multiplicity

In this appendix, we consider REE multiplicity in the context of the simple NK model with only sticky prices that we study in the main text. Some of the material in this appendix appeared in earlier versions of this paper (see Christiano et al. (2018)). We consider stability under learning (E-stability) of the different REEs. See Evans and Honkapohja (2001) for discussion of E-stability and its use in equilibrium selection. See Boneva et al. (2016) for

a discussion of multiple REE at the ZLB. The first two sub-sections of this appendix work with the model with Rotemberg-style nominal rigidities because that model has no state variable, which simplifies the analysis. Because the specification of adjustment costs with Rotemberg-style nominal rigidities is somewhat ad hoc, and because different specifications lead to the same first-order approximation to the model, we also consider multiplicity in the context of the model with Calvo-style nominal rigidities in the third sub-section. We consider that model more micro-founded and, as a result, perhaps a better model for analysis of equilibrium multiplicity. However, because of the similarities between the models with Rotemberg-style nominal rigidities and Calvo-style nominal rigidities, it is unclear that either should be regarded as strictly better.

We consider equilibrium multiplicity in the fully nonlinear version of the simple NK model. Note that analyses that linearize around the target-inflation steady state while also maintaining the nonlinearity associated with the ZLB would not identify the multiplicity we highlight.

G.1 Rotemberg-Style Nominal Rigidities

In this sub-section, we focus on the version of the model with Rotemberg-style adjustment costs. Evans et al. (2008) consider a version of this model when there is no shock to the natural rate of interest. They analyze the model's two steady states (in one steady state, $\pi = 1$, in the other $\pi = \beta$).⁴⁸ They find that only the target-inflation steady state in which $\pi = 1$ is stable under learning, and they argue that learning dynamics are affected by the zero lower bound. Here, we focus on equilibria following a shock to the natural rate of interest, as in Eggertsson and Woodford (2003). Specifically, we assume that r_t is given as in section 6.2.

With Rotemberg-style adjustment costs, the model has no state variable. We consider minimum-state-variable REEs (see McCallum (1983) for related discussion). An REE is a set of values for output, employment, inflation, and consumption, $Y^\ell, N^\ell, \pi^\ell, C^\ell$, respectively, when $r = r^\ell$. We assume that the economy reverts to the unique rational equilibrium steady state, $Y^{ss}, N^{ss}, \pi^{ss}, C^{ss}$, with $R^{ss} > 1$ when $r = r^{ss}$.⁴⁹

⁴⁸See also McCallum (2002).

⁴⁹Here we only consider equilibria in which quantities and prices are constant for a given value of r . For example, we do not consider sunspot equilibria.

The four equilibrium conditions associated with the four unknowns, $\pi^\ell, C^\ell, R^\ell, N^\ell$, are

$$1 = \frac{1}{1+r^\ell} \left[p \frac{1}{\pi^\ell} + (1-p) \frac{C^\ell}{C^{ss}} \right], \quad (89)$$

$$\begin{aligned} (\pi^\ell - 1) \pi^\ell (C^\ell + G^\ell) &= \frac{\varepsilon - 1}{\phi} (\chi N^\ell C^\ell - 1) N^\ell \\ &+ \frac{1}{1+r^\ell} p (\pi^\ell - 1) \pi^\ell (C^\ell + G^\ell), \end{aligned} \quad (90)$$

$$N^\ell = (C^\ell + G^\ell) \left(1 + \frac{\phi}{2} (\pi^\ell - 1)^2 \right), \text{ and} \quad (91)$$


$$R^\ell = \max \{1, 1 + r^{ss} + \alpha (\pi^\ell - 1)\}. \quad (92)$$

Equations (89) and (90) take into account that $\pi^{ss} = 1$. In addition, we verify and use the fact that $R^\ell = 1$. We compute C^{ss} using the steady state of the model. Equation (90) can be expressed as one equation in the unknown, π^ℓ , after using equations (89) and (91), to express C^ℓ and N^ℓ as functions of π^ℓ . Then, we can find a candidate equilibrium by finding a value of π^ℓ that sets a function, $f(\pi^\ell) = 0$. To verify that a candidate value of π^ℓ is an equilibrium, we must verify that the implied aggregate quantities and firm values are non-negative.

Our baseline parameters are the same as in the main text. In the $R > 1$ steady-state REE, $C^{ss} = 0.8$, $\pi^{ss} = 1$, $N^{ss} = 1$. While $r = r^\ell$, we set $G^\ell = G(r^\ell) = G^{ss}$. We also consider an alternative specification for government purchases, given by

$$G^\ell = G(r^\ell) = 1.05 \times G(r^{ss}). \quad (93)$$

Figure 5: $f(\pi^\ell)$ Corresponding to the Target-Inflation Steady-State Equilibrium



Note: The zeros of the function f , which is defined in the text, identify REE values for π^ℓ assuming that the economy goes to the target-inflation steady state after $r_t = r^{ss}$. The dashed line is the case when G^ℓ is high. Source: Authors' calculations.

Figure (5) displays the function $f(\pi^\ell)$ for a range of values of π^ℓ in the baseline (solid blue line) and alternative (dashed blue line) cases. In each case, there are two values of π^ℓ for which $f(\pi^\ell) = 0$. Table 5 reports the values of C^ℓ , w^ℓ , N^ℓ , R^ℓ and π^ℓ at these zeros of f . Each crossing corresponds to an interior equilibrium in which the ZLB binds.

Table 5: Equilibrium Values While $r_t = r^\ell$, Returning to Target-Inflation Steady State

	REE A	REE B		REE A'	REE B'
$400(\pi^\ell - 1)$	-35.78	-6.60	$400(\pi^\ell - 1)$	-36.99	-3.00
$400(R^\ell - 1)$	0	0	$400(R^\ell - 1)$	0	0
C^ℓ	0.48	0.74	C^ℓ	0.47	0.77
N^ℓ	0.98	0.95	N^ℓ	1.00	0.98
w^ℓ	0.59	0.88	w^ℓ	0.58	0.95
			$\frac{\Delta C + \Delta G}{\Delta G}$	-0.17	3.95

(a) $G^\ell = G^{ss}$

(b) $G^\ell = 1.05 \times G^{ss}$

Note: This table reports $\{\pi^\ell, R^\ell, C^\ell, N^\ell, w^\ell\}$ for two equilibria indicated by A and B when $G^\ell = G^{ss}$ (6a) and when $G^\ell = 1.05G^{ss}$ (6b). Each equilibrium returns to the target-inflation steady state as soon as $r = r^{ss}$. The government purchases multiplier reported in the last line of panel is the change in GDP per unit increase in G within each of the type A and B equilibria. Source: Authors' calculations.

With regard to the fiscal multiplier, we calculate the effect of an increase in G comparing A to A' and B to B' —that is, comparing two Bad-ZLB equilibria and two Good-ZLB equilibria (see Figure 5). Table 5 shows that the multiplier is very large in the latter case and very small in the former case. Consistent with this observation, expected deflation is much larger at A' than at B' .

When we apply the methodology of section E, we find that the REE associated with B is E-stable ($b < 1$) while the REE associated with A is not E-stable ($b > 1$).

G.2 Forward Guidance in the Model with Rotemberg-Style Nominal Rigidities

In this subsection, we consider multiplicity of equilibria in the context of forward guidance, as described in Section 7. We work with the simple NK model with only sticky prices and Rotemberg-style nominal rigidities. The model has a shocks to the natural rate of interest, as in the previous sub-section. The monetary authority commits to keeping the nominal interest rate at the ZLB for J periods after the discount rate has returned to its steady-state level. To make our point as simply as possible, we consider the case $J = 1$. In this sub-section we show that the number of REE proliferates under forward guidance.

In period I , $R = 1$ even though $r_t = r^{ss}$. People know that the economy will transition to steady state in period $I + 1$. The equilibrium conditions in period I are equations (89)

through (91) adjusted for forward guidance:


$$1 = \frac{1}{1 + r^{ss}} \frac{C^I}{\pi^{ss} C'^{ss}} \quad (94)$$

$$0 = (\pi^I - 1) \pi^I (C^I + G^{ss}) - \frac{\varepsilon - 1}{\phi} (\chi N^I C^I - 1) N^I \quad (95)$$

$$N^I = (C^I + G^I) \left(1 + \frac{\phi}{2} (\pi^I - 1)^2 \right). \quad (96)$$

Equations (94) and (96) define functions mapping π^I to C^I and N^I . These functions allow us to express the left-hand side of equation (95) as a function of π^I . We denote this function by $f^I(\pi^I)$. A candidate continuation equilibrium in period I is a value of π^I such that $f^I(\pi^I) = 0$ along with the associated values of C^I, N^I, w^I , and the present value of the intermediate good firm in period I . The four variables must be non-negative for a candidate equilibrium to be an equilibrium. Figure 6 displays the f^I function for a range of values of π^I . We find two continuation REE corresponding to the two zeros of f^I displayed in the figure (see points $\mathcal{A}$ and $\mathcal{B}$).⁵⁰ The REE at $\mathcal{B}$ features very high inflation, and price adjustment costs are very unrealistically high.

Figure 6: Equilibria in Period of Switch from $r_t = r^\ell$ to $r_t = r^{ss}$ Under One-Period Forward Guidance



Notes: Graph of the function, $f^I(\pi^I)$, discussed after equation (96). The two crossings with the zero line correspond to REE in period I , the date when r switches from $r_t = r^\ell$ to $r_t = r^{ss}$. Monetary policy in period I corresponds to one-period forward guidance—that is, the interest rate is held at zero in period I and then reverts to R^{ss} . The red star indicates the level of inflation in period I in the absence of forward guidance.

We now compute the REE allocations in the periods before I , which we denote by I_{-1} , conditional on the continuation equilibrium starting in period I . The period I_{-1} equilibrium

⁵⁰From equation (94) we see that C_I does not vary with π_I . It follows that f_I is quadratic function of π_I , so that the two solutions displayed in Figure 6 are the only zeros of f_I .

conditions are the appropriate analog of equations (89) through (91):

$$1 = \frac{1}{1+r^\ell} \left[p \frac{C^\ell}{\pi^\ell C^\ell} + (1-p) \frac{C^\ell}{\pi^I C^I} \right] \quad (97)$$

$$\begin{aligned} & (\pi^\ell - 1) \pi^\ell (C^\ell + G^\ell) - \frac{\varepsilon - 1}{\phi} (\chi N^\ell C^\ell - 1) N^\ell \\ & - \frac{1}{1+r_\ell} \left[p (\pi^\ell - 1) \pi^\ell (C^\ell + G^\ell) + (1-p) (\pi^I - 1) \pi^I \frac{C^\ell}{C^I} (C^I + G^{ss}) \right] = 0 \end{aligned} \quad (98)$$


$$N^\ell = (C^\ell + G^\ell) \left(1 + \frac{\phi}{2} (\pi^\ell - 1)^2 \right) \quad (99)$$

Here, we impose the condition that $R^\ell = 1$, which we verify.

Given C^I and π^I , equations (97) through (99) define a mapping from π^ℓ to C^ℓ and N^ℓ . Now, we can express the left-hand side of equation (98) as a function of π^ℓ . We denote this function by $f^{I-1}(\pi^\ell; \pi^I, C^I)$. There are two functions, f^{I-1} , conditional on the π^I, C^I associated with the period I continuation equilibria, $\mathcal{A}$ and $\mathcal{B}$.

Figure 7 displays both f^{I-1} functions for a range of values of π^ℓ ; see the dotted and dot-dashed lines. We chose the range of π^ℓ so that the graph only displays zeros of f^{I-1} that correspond to equilibria. We find two equilibria corresponding to the f^{I-1} associated with $\mathcal{A}$ (see A and B in Figure 7) and one associated with $\mathcal{B}$ (see C in Figure 7). So there are three REEs with forward guidance. The two REEs without forward guidance can be seen in the solid line in Figure 7 (we take this curve from Figure 5).

Figure 7: REE Equilibria at the ZLB with and without Forward Guidance



Notes: The solid line reproduces the solid line in Figure 5 and corresponds to the case of no forward guidance. The dashed and dot-dashed lines correspond to the case of forward guidance. The dashed line corresponds to the case in which the economy goes to point $\mathcal{B}$ in the period of the switch in r_t to r^{ss} (period I). It crosses the zero line more than once, but the other crossing involves very high inflation and is not an equilibrium because the present value of intermediate goods monopolists is negative. The dot-dashed line corresponds to the case in which the economy goes to point $\mathcal{A}$ in period I (see Figure 6).

G.3 Calvo-Style Nominal Rigidities

In this sub-section, we work with a model of Calvo-style nominal rigidities, but that is otherwise the same as in the sub-section (G.1).⁵¹ The nonlinear equations defining an REE in this model are $\frac{1}{C_t} = R_t \frac{1}{1+r_t} \mathbb{E}_t \frac{1}{\pi_{t+1} C_{t+1}}$, $F_{P,t} = \frac{C_t + G_t}{C_t} + \frac{1}{1+r_t} \xi_P \mathbb{E}_t \pi_{t+1}^\varepsilon F_{P,t+1}$, $Z_{P,t} = \chi N_t (C_t + G_t) + \frac{1}{1+r_t} \xi_P \mathbb{E}_t \pi_{t+1}^{\varepsilon-1} Z_{P,t+1}$, $p_t^* N_t = C_t + G_t$, $\tilde{p}_t = \frac{Z_{P,t}}{F_{P,t}}$, $1 = (1 - \theta_P) \tilde{p}_t^{1-\varepsilon} + \xi_P \pi_t^{\varepsilon-1}$, $(p_t^*)^{-1} = (1 - \theta_P) \tilde{p}_t^{-\varepsilon} + \xi_P \pi_t^\varepsilon (p_{t-1}^*)^{-1}$, $R_t = \max \{1, \beta^{-1} + \alpha_\pi (\pi_t - 1)\}$. When $r_t = r^{ss}$, we assume that the system converges to the target-inflation steady state. In that solution, p_{t-1}^* is the only state variable. In that solution, let $X_t = X^{ss}(p_t^*)$. Because $r_t = r^{ss}$ is an absorbing state, we can solve for these decision rules without reference to the solution when $r_t = r^\ell$.

When $r_t = r^\ell$, the nonlinear equations defining an REE in this model are the same as in the previous paragraph, but we can write the first three as


$$\begin{aligned} \frac{1}{C_t} &= R_t \frac{1}{1+r^\ell} \frac{p}{\pi_{t+1} C_{t+1}} + R_t \frac{1}{1+r^\ell} \frac{1-p}{\pi^{ss}(p_t^*) C^{ss}(p_t^*)} \\ F_{P,t} &= \frac{C_t + G_t}{C_t} + \frac{1}{1+r^\ell} \xi_P (p \pi_{t+1}^{\varepsilon-1} F_{P,t+1} + (1-p) p [\pi^{ss}]^{\varepsilon-1} (p_t^*) F_P^{ss}(p_t^*)) \\ Z_{P,t} &= \chi N_t (C_t + G_t) + \frac{1}{1+r^\ell} \xi_P (p \pi_{t+1}^\varepsilon Z_{P,t+1} + (1-p) p [\pi^{ss}]^\varepsilon (p_t^*) Z_P^{ss}(p_t^*)). \end{aligned}$$

In the context of a nonfundamental shock, Mertens and Ravn (2014) noted that this system of equations defines a nonlinear function of p_t^* . Further, they noted that if the shock remained in this state for a long time then p_{t-1}^* would converge to a constant (and so would the other variables). Here, we have a fundamental shock, r_t . However, the system still defines a nonlinear function of p_{t-1}^* if $r_t = r^\ell$ for a long time then p_{t-1}^* would converge to a constant (and so would the other variables). If we consider only the point to which the system could converge if $r_t = r^\ell$ for a long time, and if we only consider interior equilibria, these equations define a function $f^C(\pi^\ell)$ that must be zero at any limit point of an REE.

Figure (8) displays the function $f_C(\pi^\ell)$ for a range of values of π^ℓ . There are two values of π^ℓ for which $f_C(\pi^\ell) = 0$. We also consider the case when $G^\ell = 1.05 \times G^{ss}$ (the dashed lines). As was the case in the model with Rotemberg-style nominal rigidities, there are two values of π^ℓ for which $f_C(\pi^\ell) = 0$. Additionally, only the crossings on the right are E-stable.

⁵¹Unlike in the model with Rotemberg-style nominal rigidities, we set $(1+r^\ell)^{-1} = 1.005$ and $p = 0.75$.

Figure 8: $f^C(\pi^\ell)$ Corresponding to the Target-Inflation Steady-State Equilibrium



Note: The function, f_C , is defined in the text. The dashed line is the case when G_t is high. Source: Authors' calculations.

H The Learning Equilibrium Studied by Mertens and Ravn (2014)

In this appendix, we explore the findings of Mertens and Ravn (2014) using the simple NK model with only sticky prices and parameter values that we study in our paper. Mertens and Ravn (2014) also study a simple NK model with Calvo-style nominal rigidities and only sticky prices. They show that there exists a nonfundamental REE. In that equilibrium, the natural rate of interest is always equal to its steady state value. Similar to the stochastic process for r_t in section 6.2, the nonfundamental shock, S_t , follows a stochastic process to that

$$\begin{aligned} \Pr[S_{t+1} = 1|S_t = 1] &= p, \quad \Pr[S_{t+1} = 0|S_t = 1] = 1 - p, \\ \Pr[S_{t+1} = 1|S_t = 0] &= 0. \end{aligned} \tag{100}$$

Notice that $S_t = 0$ is an absorbing state. When $S_t = 1$, Mertens and Ravn (2014) show that there is an REE where $R_t = 1$ while $S_t = 1$. In this section, we set $p = 0.9$.

H.1 REE

When $S_t = 1$, the nonlinear equations defining an REE in this model are the same as in Appendix G.3, but we can write the first three as

$$\begin{aligned}\frac{1}{C_t} &= R_t \beta \frac{p}{\pi_{t+1} C_{t+1}} + R_t \beta \frac{1-p}{\pi_{ss}(p_t^*) C_{ss}(p_t^*)} \\ F_{P,t} &= \frac{C_t + G_t}{C_t} + \beta \xi_P \left(p \pi_{t+1}^{\varepsilon-1} F_{P,t+1} + (1-p) p [\pi_{ss}]^{\varepsilon-1} (p_t^*) F_P^{ss}(p_t^*) \right) \\ Z_{P,t} &= \chi N_t (C_t + G_t) + \beta \xi_P \left(p \pi_{t+1}^{\varepsilon} Z_{P,t+1} + (1-p) p [\pi_{ss}]^{\varepsilon} (p_t^*) Z_P^{ss}(p_t^*) \right).\end{aligned}$$

Mertens and Ravn (2014) noted that this system of equations defines a nonlinear function of p_t^* . Further, they noted that if $S_t = 1$ for a long time then p_{t-1}^* would converge to a constant (and so would the other variables).

H.2 Euler-Equation Learning

Mertens and Ravn (2014) consider a learning equilibrium using the Euler equation approach to learning (see Appendix C for a discussion of Euler equation learning). To implement the Euler equation approach with the zero lower bound, Mertens and Ravn (2014) log-linearize the equilibrium conditions around the target-inflation steady state but preserve the nonlinearity in the monetary policy rule. Mertens and Ravn (2014) assume that agents know that $\hat{C}_t = \hat{\pi}_t = 0$ when $S_t = 0$. When $S_t = 1$, the system of equations can be written as

$$\begin{aligned}\hat{C}_t &= p \mathbb{E}_t \hat{C}_{t+1} + p \mathbb{E}_t \hat{\pi}_{t+1} - \max \{ \beta \log(\beta), \beta \alpha_{\pi} \hat{\pi}_t \} \\ \hat{\pi}_t &= \frac{(1 - \beta \xi_P)(1 - \xi_P)}{\xi_P} \left[(1 + C) \hat{C}_t + G \hat{G}_t \right] + \beta p \mathbb{E}_t \hat{\pi}_{t+1}\end{aligned}$$

In a learning equilibrium, beliefs are set so that $\mathbb{E}_t \hat{C}_{t+1} = m_{C,t-1}$ and $\mathbb{E}_t \hat{\pi}_{t+1} = m_{\pi,t-1}$. These values are then updated according to

$$m_{i,t} = m_{i,t-1} + \gamma_t \left(\hat{i}_t - m_{i,t-1} \right).$$


Mertens and Ravn (2014) use constant gain learning ($\gamma_t = 0.1$). As in the rest of our paper, we use decreasing gain and set $\gamma_t = \frac{1}{t+1}$.

As in Mertens and Ravn (2014) we define the government purchases multiplier to be $\frac{\Delta C_t + \Delta G_t}{\Delta G_t}$ where Δ indicates the difference relative to the case when $G_t = 0$. Mertens and Ravn (2014) consider two experiments. In the first experiment, Mertens and Ravn (2014) calculate the government purchases multiplier in the REE. In the second experiment, Mertens and Ravn (2014) calculate the size of the government purchases multiplier under learning. To calculate the multiplier under learning, Mertens and Ravn (2014) set $\hat{G}_t = 0$ and initial

beliefs about consumption and inflation to be somewhat above their REE values. They then calculate a path for C_t . Then Mertens and Ravn (2014) set $\hat{G}_t = \hat{G} > 0$ and also change initial beliefs to be lower, similar to the way initial beliefs are lower in the REE when $\hat{G}_t = \hat{G} > 0$. They then calculate an alternative path for C_t and compute the multiplier.

We consider a third experiment in which $\hat{G}_t = \hat{G} > 0$ and initial beliefs are set to the same initial value as in the learning equilibrium when $\hat{G}_t = 0$. Figure 9a displays the value of the government purchases multiplier in these three experiments under the assumption that $S_t = 1$ for the entire duration of the period shown. Note that the REE multiplier is low. As in Mertens and Ravn (2014), we find that the compound, second experiment in which beliefs and G_t change exhibits a low multiplier. In the third experiment, when we do not change initial beliefs for the higher value of G_t , the multiplier is larger than unity. Figure 9a is also consistent with the findings of Mertens and Ravn (2014) that Euler equation learning is slow. The value of b for the system of equations associated with Euler equation learning is 1.03, which is consistent with expectations diverging from the nonfundamental equilibrium (as reported by Mertens and Ravn (2014)), but at a slow rate.

Figure 9: Size of Government Purchases Multiplier



Note: The figure shows the size of the government purchases multiplier under three experiments in our simple NK model with only sticky prices. Experiment 1 is the REE. Experiment 2 changes initial expectations when G_t changes, as in Mertens and Ravn (2014). Experiment 3 leave initial expectations unchanged when G_t changes. Panel (a) is constructed under Euler equation learning, as in Mertens and Ravn (2014). Panel (b) is constructed under anticipated utility as outline in Appendix E. The figure displays the value of the government purchases multiplier in learning equilibria and the REE under the assumption that $S_t = 1$ for the entire duration of the period shown. This is a low probability event that would be surprising. Source: Authors' calculations.

H.3 Anticipated Utility

Straightforward modifications of the approach outlined in Appendix E allow us to solve for the linear approximation to the learning equilibrium under anticipated utility using the methodology of Appendix B. Note that we solve for the limit point of p_t^* (and the other variables) while $S_t = 1$ linearize the equations around this limit point while $S_t = 1$.⁵² We also linearize the equations once $S_t = 0$ around the target inflation steady state.

We conduct the same three experiments that we described in the previous section. There are dynamics associated with the REE in Figure 9b because p_t^* is a state variable of the model and it is linearized around a value different from unity. The value of the multiplier in the REE is somewhat different from the value in Figure 9a because we are linearizing around the nonlinear REE, rather than around the target inflation steady state. As in the previous sub-section, we find that the compound, second experiment in which beliefs and G_t change exhibits a low multiplier but the multiplier is larger when we do not change initial beliefs (experiment 3).

It is notable that the dynamics in Figure 9b are also consistent with slow learning under anticipated utility in the compound second experiment. The value of b under anticipated utility is 1.7. Because $b > 1$, the system diverges from the sunspot REE that Mertens and Ravn (2014) study, however, as discussed above, this happens slowly for beliefs that are close enough to the REE in the compound second experiment. For experiment 3 (when we do not change beliefs at the same time that we change G_t at the ZLB), divergence from the REE happens quickly because initial beliefs when $\hat{G}_t > 0$ are not so close to the REE. In this case, we only show the first few quarters of the simulation because the ZLB would no longer be binding after that period in the simulation and the linearization of the model would then have to incorporate the interest rate response from the Taylor rule.

⁵²It turns out that the linear approximation to the REE when $S_t = 1$ is locally indeterminate. We resolve the indeterminacy by requiring that in the REE all variables converge at a rate faster than 0.90^t . This results in a minimum state variable linear approximation to the REE, which we utilize to construct our learning equilibrium as in appendix B. See McCallum (1983) for a discussion regarding reasons to select the minimum state variable solution.