

Slow Learning

Lawrence J. Christiano Martin Eichenbaum
Benjamin K. Johannsen*

January 26, 2025

Abstract

This paper analytically characterizes the speed of convergence under learning to a rational expectations equilibrium (REE) for a large class of multivariate models in which people's beliefs about model outcomes are central determinants of those outcomes. The paper then investigates what features of an economy determine whether convergence under learning is fast or slow. We do so by applying our analytic results to the New Keynesian model and studying the impact of the Zero Lower Bound (ZLB) on the speed of convergence to a REE. Under certain circumstances, convergence of a learning equilibrium to the REE equilibrium can be so slow that policy analysis based on rational expectations is very misleading.

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

1 Introduction

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

To address this question, we analytically characterize the speed of convergence in a broad class of non-linear, non-stochastic learning models. That class of models has two important characteristics. First, people learn using either standard Bayesian methods or least-squares. Second, people’s beliefs about model outcomes are central determinants of equilibrium outcomes.

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

Our central finding is that when beliefs are partially self-fulfilling, learning equilibria converge slowly to rational expectations. Indeed, learning can be extraordinarily slow, with progress being measured in millennia. Under these circumstances, policy analyses based on rational expectations can be very misleading. As Vives (1993) writes, in a changing world, for all practical purposes, “‘slow’ convergence may mean no convergence.”

We apply our propositions to analyze speed of convergence of learning equilibria in the non-linear new-Keynesian model when the zero lower bound (ZLB) is and is not binding.² We argue that when the ZLB is binding, the crucial parameter b is close

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

²Much of the work in the initial aftermath of that event combined rational expectations with the NK model. See, for example, Eggertsson and Woodford (2004), Christiano et al. (2011) and Del Negro

to 1, so learning is *very* slow. It is so slow, that the ZLB would be over well before the learning and rational expectations equilibria are reasonably similar. Moreover, the effects of monetary and fiscal policies are very different in the two equilibria. In contrast, when the ZLB is not binding, b is much smaller, and the learning model converges rapidly to an REE.

Convergence is slow when the ZLB is binding because the expectations of households and firms are partially self-fulfilling. To understand why learning is slow, suppose that firms and households expect lower inflation in the future. Because of price-setting adjustment costs, firms are incentivized to cut prices today. In the ZLB, low inflation expectations mean households believe the real interest rate is high. Consequently, households reduce their demand for consumption, which leads to a fall in the marginal cost of production. So, the actions of both households and firms lead to lower current inflation, consistent with their initial beliefs. With learning, low current inflation shifts expected inflation down in the next period. The previous mechanism repeats itself in the next period so that actual inflation in the next period is also low. We conclude that, in the ZLB, deflation expectations are partially self-fulfilling, and the NK model behaves like a high b economy. In sharp contrast, when the ZLB is not binding and monetary policy is governed by a Taylor rule, people's expectations about inflation are not self-fulfilling, and the parameter b is small.³

The speed of convergence plays a crucial role in analyzing both the effects of shocks to the economy and the efficacy of various policies in dealing with those shocks. In the NK model, when people have rational expectations, a shock that triggers a binding ZLB leads to a sharp decline in inflation and output (see Eggertsson and Woodford (2004)). The large effects arise because the shock triggers low expected inflation and high real interest rates. Under learning the same shock leads only to a moderate and gradual decline in inflation and a moderate rise in real interest rates because expectations are partially backward-looking.

Turning to fiscal policy, we find that the efficacy of a rise in government purchases is much smaller under learning than under rational expectations. Under rational expectations, the multiplier is very large in the ZLB because an increase in government purchases causes a rise in 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

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

in consumption, and a multiplier substantially larger than unity. Under learning, expected inflation is partially backward-looking and doesn't move much after an increase in government purchases. So, the real interest rate doesn't fall by very much, the key driver of the large REE multiplier is effectively eliminated, and the multiplier is close to unity.

We also analyze the effects of forward guidance when the ZLB is binding. To simplify the analysis, 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. We show that the number of REEs proliferates under forward guidance, but only one REE is stable under learning. Consistent with the existing literature (for example, Del Negro et al. (2023) and Woodford (2012)), we find that even this simple form of forward guidance is powerful under rational expectations. As is well-known, the power of forward guidance under rational expectations reflects its strong effect on expected inflation. Under learning, expectations are partially backward-looking, and forward guidance is not very powerful. So, as with fiscal policy, a REE-based analysis of monetary policy can be very misleading.

Our propositions refer to the asymptotic rate of convergence. A natural question is whether the propositions are useful for characterizing convergence speed over short horizons. Based on our analysis of the NK model, we show that the answer is yes.

In our analysis, people fully integrate the fact that they are learning when they solve their problems, that is, households and firms are internally rational in the sense defined by Adam and Marcet (2011). Implementing internal rationality in the non-linear solution of the model is computationally very challenging.⁴ The reason is that the parameters characterizing beliefs are state variables, an effect that greatly exacerbates the curse of dimensionality. The associated computational burden explains why much of the learning literature works with a version of Kreps (1998)'s *Anticipated Utility* approach. In this approach, people update their beliefs every period as new data come in. But, when they make their decisions, people proceed as though their beliefs will never be revised again.⁵

We simulate our model using internal rationality and anticipated utility. We find

⁴We solve our model using a compiled programming language (c++) and we make use of more than 300 processors. See Appendix B for details.

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

that the results are qualitatively similar. However, for some experiments, there are important quantitative differences between the two approaches due to the more prominent role played by uncertainty under internal rationality. These results are consistent with those obtained by Cogley and Sargent (2008), who studied a stochastic endowment economy with a storage technology.

The remainder of this paper is organized as follows. Section 2 discusses related literature. Section 3 states the propositions, which characterize the asymptotic rate of convergence of a learning equilibrium for a broad class of models (proofs are provided in the appendix). Section 4 lays out the non-linear NK model under learning and rational expectations. Section 5 analyzes multiplicity of REE in the non-linear NK model. Section 6 analyzes the local and global learnability of those equilibria. In Section 7, we analyze the speed of convergence of learning equilibria in the non-linear NK model. Section 8 assesses the sensitivity of the efficacy of fiscal policy and forward guidance to learning. Section 9 establishes the value of b in the NK model. Section 10 contains concluding remarks.

2 Related Literature

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

The second is the literature that studies the properties of recursive stochastic estimators in learning models. Ljung (1977) establishes that a recursive estimator, $\hat{\theta}_t$, converges almost surely to a limiting value, θ , if a particular ordinary differential equation (ODE), determined by the economic model, has eigenvalues with real parts that are less than unity. Marcet and Sargent (1989b; 1989a), Woodford (1990), Evans and Honkapohja (2000; 2001), and others build on Ljung (1977) to study the conditions under which learning equilibria converge to an REE. Marcet and Sargent (1995) numerically study the rate at which these learning equilibria converge to an REE. Christopheit and Massmann (2018) provide analytic results in a linear, scalar stochastic model. In contrast, we provide multivariate, non-stochastic results and apply them to the non-linear NK model.

Ferrero (2007) discusses learning in the context of a linear NK model in which the

ZLB on interest rates is not binding. He uses the simulation methods proposed by Marcet and Sargent (1995) to study convergence rates of learning equilibria. Ferrero (2007) adopts the so-called Euler-equation approach to learning as opposed to our approach; see Evans (2021) for a definition of the Euler-equation approach to learning, and see Preston (2005) and Adam and Marcet (2011) for a critique of that approach.⁶ Another difference with Ferrero (2007) is that we compare convergence rates in a non-linear NK model when the ZLB on interest rates is and is not binding.

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

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

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

Heemeijer et al. (2009) and Hommes (2011) study positive and negative feedback loops from expectations to outcomes using laboratory experiments and univariate models with constant gain. In contrast, we analytically characterize convergence rates of beliefs for a broad class of models and numerically analyze rates of convergence in a multivariate NK model under Bayesian learning.

Our paper is also related to a recent game-theoretic grounded literature that analyzes the implications of bounded rationality for the effectiveness of fiscal and monetary policy. Farhi and Werning (2019) use k -level thinking models to study how deviations

⁶Our approach is an example of what Evans (2021) calls the agent-based approach to learning.

from rational expectations affect the effectiveness of forward guidance. García-Schmidt and Woodford (2019) study forward guidance and interest rate pegs using reflective expectations. Iovino and Sergeyev (2023) apply k -level thinking and reflective expectations to analyze the effects of quantitative easing. Angeletos and Lian (2017) develop the idea that a lack of common knowledge can attenuate general-equilibrium effects and damp the effects of government spending. Angeletos and Lian (2017; 2018) analyze the consequences of bounded rationality for the size of fiscal multipliers.

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

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

3 Learning in a Non-Linear Environment

Here, we consider the speed of convergence of learning in the following non-linear environment, which is very similar to that studied in Evans and Honkapohja (2000). Let θ_t be a k -dimensional vector of variables which summarizes people's period t priors about a set of variables that will be determined at time $t + 1$. We interpret θ_t as a deviation from a particular fixed point of beliefs in our learning algorithm.⁷ Given a set of beliefs, θ_{t-1} , the environment generates outcomes in period t according to the non-linear function, $M(\theta_{t-1}, \gamma_t)$. 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 $t = 1, 2, 3, \dots$. Here, θ_0 is given and $\gamma_t = \frac{1}{c_1+t}$ for $c_1 \geq 0$ is the gain. This type of gain parameter emerges from standard Bayesian learning as well as least squares learning.

⁷We do not require that this fixed point be unique.

The vector-valued function, $M : \mathbb{R}^k \times \mathbb{R} \rightarrow \mathbb{R}^k$ has the following properties: $M(0, 0) = 0$; M is continuously differentiable in a neighborhood of the origin; $M(0, \gamma_t) = 0$; and $M(0, \gamma_t)$ is continuously differentiable in a neighborhood of $(0, \gamma_t)$. Let D_1M denote the derivative of M with respect to the vector θ_{t-1} (which means it is a $k \times k$ matrix). We assume the real parts of the eigenvalues of D_1M are strictly less than unity. The scalar, b , denotes the largest real part of the eigenvalues of D_1M .

Evans and Honkapohja (2000) consider the same environment as us, but with a more general specification of γ_t .⁸ They establish that for, (i) $b < 1$, and (ii) θ_0 sufficiently close to 0, θ_t converges to zero. Christopeit and Massmann (2018) consider a scalar, linear, stochastic version of our environment. Their results imply that for the non-stochastic environment studied here, θ_t evolves as t^{b-1} . We extend this result to a multivariate, non-linear environment.⁹

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

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

Here, $\|\cdot\|$ denotes a norm on \mathbb{R}^k . The first part of definition 1 says that for any positive δ , $\|x_t\|$ asymptotically converges no slower than $t^{b-1+\delta}$. The second part says that $\|x_t\|$ asymptotically converges no faster than $t^{b-1-\delta}$. In this sense, b characterizes the power rate of convergence. Importantly, two series that asymptotically converge to zero at the rate t^{b-1} may behave very differently for finite T . Some examples of series that are different even for large t even though they asymptotically converge to zero at the rate t^{b-1} , include $x_t = \log(t) t^{b-1}$, $y_t = t^{b-1}/\log(t)$, and $z_t = [2 + \sin(t)] t^{b-1}$. These series asymptotically converge at the same rate when considering only rates of power convergence, which is what our definition captures. There are, of course, series that converge at a faster rate than power convergence, such as those that converge at a geometric rate (for example, $x_t = \rho^t$ for $|\rho| < 1$). As it turns out, the NK model with

⁸They assume $1 > \gamma_t > 0$, $\gamma_t \rightarrow 0$ and $\lim_{T \rightarrow \infty} \sum_{t=1}^T \gamma_t$ diverges. Our specification of γ_t satisfies these conditions.

⁹Ljung (1977) showed that to determine whether θ_t converges, it is useful to consider the ordinary differential equation, $\dot{\theta}(\tau) = D_1M\theta(\tau) - \theta(\tau)$, where τ evolves in continuous time. In the scalar case, $b = D_1M$ and the solution to this equation is $\theta(\tau) = [e^{(b-1)\tau}]^T \theta(0)$. The same parameter, b , determines whether $\theta(\tau)$ and θ_t converge as well as their rates of convergence. But, the rates of convergence of these variables are qualitatively different: $\theta(\tau)$ converges at a geometric rate in τ -time while θ_t converges at a power rate in t -time. See the appendix for further discussion in the context of a linear model.

Bayesian learning (or any model whose reduced form satisfies equation 1) rate exhibits power convergence.

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

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

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

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

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

4 Learning in the New Keynesian Model

In this section, we describe a simple NK model that has been widely used to study the effects of the ZLB. As in Eggertsson and Woodford (2003), we allow for a shock to the household's discount rate that can cause the ZLB on the interest rate to be binding. It is convenient to express people's problems in recursive form.

In the current period, households discount next period's utility by $1/(1+r)$. In steady state, $r = r_{ss} > 0$. We assume that initially the economy is in the unique non-stochastic rational expectations steady state in which the nominal interest rate is positive. Then, unexpectedly, $r = r_\ell < r_{ss}$. People correctly understand that the next period's discount rate, r' , is drawn from a two-state Markov chain, $r' \in [r_\ell, r_{ss}]$, with

an absorbing state:

$$\begin{aligned} \Pr[r' = r_\ell | r = r_\ell] &= p, & \Pr[r' = r_{ss} | r' = r_\ell] &= 1 - p, \\ \Pr[r' = r_\ell | r = r_{ss}] &= 0. \end{aligned} \tag{2}$$

Once $r = r_{ss}$, the economy returns to the initial rational expectations steady state. There is another rational expectations steady state in which there is deflation and the nominal interest rate is unity (see Benhabib et al. (2001)). We abstract from that steady state equilibrium because it is not stable under the learning models that we consider.¹⁰ Moreover, focusing on one steady state greatly simplifies our analysis.

4.1 Monetary and Fiscal Policy

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

$$R = \max \{1, 1 + r_{ss} + \alpha(\pi - 1)\}, \tag{3}$$

where π is the gross rate of consumer price inflation, $\alpha/(1 + r_{ss}) > 1$, and the max operator reflects the ZLB constraint. Later, we discuss other variations on monetary policy, including forward guidance.

Fiscal policy sets the level of real government purchases, G . We consider two specifications for G . In the baseline specification, G is equal to its non-stochastic steady-state value, G_{SS} . We also consider a policy where $G = G_\ell > G_{ss}$ while $r = r_\ell$. The government also pays a wage subsidy to intermediate goods firms, which it picks to offset steady state monopoly distortions. The government finances its expenditures with lump-sum taxes and balances its budget in each period.

4.2 Private Agents' Problems

Below, we define the household and firm problems.

4.2.1 The Household's Problem When $r = r_\ell$

The household enters a period with a stock of bonds, $b_h = B_{h,t-1}/P_{t-1}$. Here, $B_{h,t-1}$ denotes the beginning-of-period t payoff on nominal bonds acquired in the previous

¹⁰See Arifovic et al. (2018) for a discussion of stability for other learning models in which the deflationary steady state is learnable.

period, when the price of consumption goods was P_{t-1} . At the beginning of a period, before markets open, the household also knows the value the vector, Θ , which summarizes its beliefs about the distribution of a vector, x :

$$x = \begin{bmatrix} C \\ \pi \end{bmatrix}.$$

Here, C and π denote the current period's aggregate consumption and aggregate inflation. The variable, π , corresponds to P_t/P_{t-1} , where P_t and P_{t-1} denote the current and previous period's aggregate price level, respectively.

In a standard recursive equilibrium, people know current-period market prices and profits when they make their current decisions. Typically, when markets open in these models, people can deduce the prices and profits from a small set of variables. In our context, these variables are the two components of x . In this spirit, we assume that people observe x when markets open, and they make their current consumption, saving, and labor decisions. In making those decisions, households internalize the effect of x on their beliefs about the distribution x' —that is, the value of x in the next period. Those beliefs, Θ' , are given by

$$\Theta' = L(\Theta, x). \quad (4)$$

The form of L depends on the model of learning being analyzed. The household takes into account uncertainty about the distribution of x and the fact that beliefs about that distribution will evolve as new data arrive (see Section 4.3.2). That is, the household is internally rational in the sense of Adam and Marcet (2011).

Let C_h, N_h, b'_h denote the representative household's consumption, hours worked and end-of-period bond holdings. The household solves

$$\max_{C_h, N_h, b'_h} \left\{ \log(C_h) - \frac{\chi}{2} (N_h)^2 + \frac{1}{1+r_\ell} [(1-p)V_{h,ss}(b'_h) + p\mathbb{E}_{\Theta'} V_h(b'_h, \Theta', x')] \right\} \quad (5)$$

subject to

$$C_h + \frac{b'_h}{R(x)} \leq \frac{b_h}{\pi(x)} + w(x)N_h + T(x). \quad (6)$$

Here, $T(x)$ denotes profits net of lump-sum taxes, $w(x)$ denotes the real wage, $R(x)$

denotes the nominal rate of interest, and $\pi(x)$ denotes the inflation rate.¹¹ In equation (5), $V_{h,ss}(b'_h)$ denotes the value function of the household conditional on $r' = r_{ss}$, and $V_h(b'_h, \Theta', x')$ denotes the value conditional on $r' = r_\ell$. The expectation operator, $\mathbb{E}_{\Theta'}$, is evaluated using the marginal data density for x' implied by $\Theta' = L(\Theta, x)$ and $r' = r_\ell$. Using the first-order optimality condition for N_h and equation (6), we reduce the household problem to finding an optimal decision rule, $b'_h(b_h, \Theta, x)$, for bond holdings.

The function, $V_{h,ss}(b_h)$, satisfies the following fixed point:

$$V_{h,ss}(b_h) = \max_{C_h, N_h, b'_h} \left\{ \log(C_h) - \frac{\chi}{2} (N_h)^2 + \frac{1}{1 + r_{ss}} V_{h,ss}(b'_h) \right\}, \quad (7)$$

subject to

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

where T_{ss} denotes steady-state profits, net of taxes, in steady state, w_{ss} denotes the steady-state real wage, R_{ss} denotes the steady-state nominal interest rate, and π_{ss} denotes the steady-state inflation rate.

The function, V_h , in equation (5) has the fixed point property:

$$V_h(b_h, \Theta, x) = \max_{C_h, N_h, b'_h} \left\{ \log(C_h) - \frac{\chi}{2} (N_h)^2 + \frac{1}{1 + r_\ell} [(1 - p) V_{h,ss}(b'_h) + p \mathbb{E}_{\Theta'} V_h(b'_h, \Theta', x')] \right\}, \quad (8)$$

where the maximization is subject to equation (6) and the law of motion for Θ in equation (4).

4.2.2 The Firm's Problem When $r = r_\ell$

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

$$Y = \left(\int_0^1 Y_f^{\frac{\varepsilon-1}{\varepsilon}} df \right)^{\frac{\varepsilon}{\varepsilon-1}}, \quad (9)$$

¹¹We constrain the choice of b'_h to a compact set $[\underline{b}, \bar{b}]$, which we discuss in Appendix B.

where $\varepsilon > 1$. The representative firm chooses inputs, Y_f , to maximize profits $YP - \int_0^1 Y_f P_f df$, subject to (9). The firm's first-order condition for the f^{th} input is

$$Y_f = \left(\frac{P_f}{P} \right)^{-\varepsilon} Y. \quad (10)$$

The f^{th} intermediate good is produced by a monopolist with production technology $Y_f = N_f$, where N_f is labor hired by firm f . Let p_f denote the f^{th} firm's price in the previous period, scaled by that period's aggregate price index—that is, $P_{f,t-1}/P_{t-1}$. Also, let p'_f denote the firm's current choice of price scaled by the current aggregate price index. In our scaled notation,

$$\frac{p'_f}{p_f} \pi = \frac{P_{f,t}}{P_{f,t-1}}. \quad (11)$$

Firms value a unit of real profits by the marginal utility of consumption, $1/C$. Prices are sticky as in Rotemberg (1982). When $r = r_\ell$ the current-period problem of firm f is to set its price p'_f so that

$$\begin{aligned} p'_f(p_f, \Theta, x) = \operatorname{argmax}_{p'_f} & \frac{1}{C(x)} \left\{ (p'_f - (1 - \nu)w(x)) (p'_f)^{-\varepsilon} Y(x) \right. \\ & \left. - \frac{\phi}{2} \left(\frac{p'_f}{p_f} \pi(x) - 1 \right)^2 (C(x) + G(r_\ell)) \right\} \\ & + \frac{1}{1 + r_\ell} [(1 - p)V_{f,ss}(p'_f) + p\mathbb{E}_{\Theta'} V_f(p'_f, \Theta', x')]. \end{aligned} \quad (12)$$

Here, $V_{f,ss}(p'_f)$ denotes the value of the firm's problem conditional on $r' = r_{ss}$ and $V_{f,ss}(p'_f, \Theta', x')$ denotes its value conditional on $r' = r_\ell$.¹² Firms and households have the same information sets and update priors in the same way. Thus, the expectations operator is the same as the one in the household's problem. In equation (12), we follow the literature by scaling price adjustment costs by real GDP.¹³ Also, ν is the government's tax subsidy on employment designed to eliminate the effect of monopoly distortions in steady state.¹⁴

¹²We constrain the choice of $\log(p'_f)$ to a compact set $[p, \bar{p}]$. See Appendix B for a discussion.

¹³See, for example, Kaplan and Violante (2018, page 711).

¹⁴That is, $(1 - \nu)\varepsilon / (\varepsilon - 1) = 1$.

The function, $V_{f,ss}(p_f)$, has the fixed-point property

$$V_{f,ss}(p_f) = \max_{p'_f} \left\{ \frac{1}{C_{ss}} \left((p'_f - (1 - \nu) w_{ss}) (p'_f)^{-\varepsilon} Y(x) \right) - \frac{1}{C_{ss}} \frac{\phi}{2} \left(\frac{p'_f}{p_f} \pi_{ss} - 1 \right)^2 (C_{ss} + G_{ss}) + \frac{1}{1 + r_{ss}} V_{f,ss}(p'_f) \right\}. \quad (13)$$

The function, V_f , in equation (12) has the fixed point property

$$V_f(p_f, \Theta, x) = \max_{p'_f} \left\{ \frac{1}{C(x)} (p'_f - s) (p'_f)^{-\varepsilon} Y(x) - \frac{1}{C(x)} \frac{\phi}{2} \left(\frac{p'_f}{p_f} \pi(x) - 1 \right)^2 (C(x) + G(r_\ell)) + \frac{1}{1 + r_\ell} [(1 - p) V_{f,ss}(p'_f) + p \mathbb{E}_{\Theta'} V_f(p'_f, \Theta', x')] \right\}. \quad (14)$$

The maximization takes into account the law of motion of Θ , controlled by L , in equation (4).

4.2.3 The Mapping from x to Aggregate Variables

For individual households' and firms' problems to be well defined, they must know the values of seven aggregate variables, $[C \ \pi \ R \ Y \ N \ w \ T]$. We assume that each agent knows the model's static equilibrium conditions so they can deduce those variables from $x = [C \ \pi]$. We denote this mapping by $F(x)$. Households derive R from π using equation (3). The mappings from x and r to Y , N , and w are given by

$$Y = (C + G(r)) \left(1 + \frac{\phi}{2} (\pi - 1)^2 \right), \quad N = Y, \quad w = \chi NC.$$

The first two equalities correspond to goods market clearing and the aggregate production function. The third equality corresponds to the belief that the labor supply curve of the individual household holds as an aggregate condition. These equalities hold in every period of our learning equilibria (described in the next sub-section).

Aggregate firm profits net of taxes implied by x and r are

$$T = (1 - w)Y - \frac{\phi}{2} (\pi - 1)^2 (C + G(r)) - G(r).$$

4.3 Equilibrium and Beliefs

The equilibrium for our model is a *learning equilibrium* for the duration of time that $r = r_\ell$, followed by a jump to the positive interest rate, steady state REE. The learning equilibrium is a sequence of period equilibria.

4.3.1 Equilibrium Definitions

We now define a period equilibrium.

Definition 2. Given Θ and r_ℓ , a period equilibrium is a set of values of x and $\Theta' = L(\Theta, x)$ such that

- (i) households and firms solve their optimization problems, defined in equations (5) and (12), respectively
- (ii) labor, goods and bond markets clear
- (iii) $p'_f = 1$, $C_h = C$, $N_h = N$

Because firms are identical, in a learning equilibrium, no firm will ever inherit a $p_f \neq 1$. Then, equation (11) and the first part of condition (iii) imply that people's views about inflation, π , are correct. The second and third parts of condition (iii) imply that people's views about C and N are correct.

The only new conditions in Definition 2 relative to those imposed by $F(x)$ are that bond markets clear ($b'_h = 0$), and firms choose $p'_f = 1$. These two conditions determine the two elements of x .

Two comments about the period equilibrium are worth emphasizing. First, people have perfect foresight regarding current aggregate variables. Second, they generally do not have perfect foresight about future aggregates. It follows that the period equilibrium under learning is generally different from what it would be if people had rational expectations.

We now define a learning equilibrium.

Definition 3. A *learning equilibrium* is :

- (i) a sequence of period equilibria in which beliefs are updated according to equation (4) when $r = r_\ell$,
- (ii) a steady state REE with $R > 1$, when $r = r_{ss}$.

In a learning equilibrium, the value of Θ in the first period when $r = r_\ell$ is exogenous. We assume that in the case of an unprecedented event, people's priors about the

economic variables, x , are very diffuse. Below, we describe how our parameterization of the initial Θ captures this property.

4.3.2 Beliefs and Equilibrium

We now describe how households' and firms' common beliefs evolve, starting in the first period that $r = r_\ell$. People assume that each of the two elements of $\log(x)$ is drawn from a Normal distribution:

$$\log(x) = \begin{bmatrix} \log(C) \\ \log(\pi) \end{bmatrix} = \begin{bmatrix} \mu_C \\ \mu_\pi \end{bmatrix} + \begin{bmatrix} \varepsilon_C \\ \varepsilon_\pi \end{bmatrix}, \quad (15)$$

$E\varepsilon_C = E\varepsilon_\pi = 0$, $E\varepsilon_C^2 = \sigma_C^2$ and $E\varepsilon_\pi^2 = \sigma_\pi^2$. These distributions are independent across time and the elements of $\log(x)$. People are uncertain about the values of μ_i , σ_i^2 for $i \in \{C, \pi\}$. 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. The prior for σ_i^2 is not exactly an inverse-gamma distribution because we truncate the support of σ_i^2 so that $E[C]$ and $E[\pi]$ have finite values. 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 parameters of the priors in the vector Θ :

$$\Theta = \left(m_C \quad m_\pi \quad 1/\lambda_C \quad 1/\lambda_\pi \quad \psi_C \quad \psi_\pi \quad 1/\alpha_C \quad 1/\alpha_\pi \right). \quad (16)$$

The posterior distribution is also proportional to the Normal inverse-gamma distribution, and the function, L , in equation (4) can be constructed using standard updating formulas, which are detailed in Appendix B.

4.3.3 Anticipated Utility

Virtually all the related literature works with a version of Kreps' *Anticipated Utility* approach to how people integrate learning into their decisions. While this approach has computational advantages, it has been criticized for being internally inconsistent (see Cogley and Sargent (2008) and Adam and Marcet (2011)). We assess the robustness of our results to using the anticipated utility approach. In our context, that

approach assumes that when households and firms make their state- x contingent decisions, they assume that in the current and all future periods, $\log(x)$ will be drawn from a Normal distribution with mean and variance *fixed* at the values of m_i and ψ_i^2 from the beginning-of-period Θ . We make two changes to the household and firm decision problems to implement this assumption. First, we set $\Theta' = \Theta$ in their next-period value functions. Second, in evaluating the expectation operator, $\mathbb{E}_{\Theta'}$, that appears in the household and firm problems, we use the log Normal density for x with mean and variance fixed at the values of m_i and ψ_i^2 from Θ . Importantly, at the beginning of the next period, firms and households set $\Theta' = L(\Theta, x)$.

In sum, anticipated utility differs from internalized learning in two ways. First, in making their state- x contingent decisions, people ignore that after they see current x , they will update their views, using $\Theta' = L(\Theta, x)$. Second, they ignore their uncertainty about the mean and variance of the distribution of $\log(x)$, and the fact that they will learn from future realizations of x .

5 Multiple Rational Expectations Equilibria

In this section, we describe the minimum state variable equilibria in our model when agents have rational expectations.

An equilibrium 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 $r = r_{ss}$.¹⁵

The four equilibrium conditions associated with the four unknowns, $\pi_\ell, C_\ell, R_\ell, N_\ell$,

¹⁵Throughout the paper, 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.

are

$$1 = \frac{1}{1 + r_\ell} \left[p \frac{1}{\pi_\ell} + (1 - p) \frac{C_\ell}{C_{ss}} \right], \quad (17)$$

$$(\pi_\ell - 1) \pi_\ell (C_\ell + G_\ell) = \frac{\varepsilon - 1}{\phi} (\chi N_\ell C_\ell - 1) N_\ell + \frac{1}{1 + r_\ell} p (\pi_\ell - 1) \pi_\ell (C_\ell + G_\ell), \quad (18)$$

$$N_\ell = (C_\ell + G_\ell) \left(1 + \frac{\phi}{2} (\pi_\ell - 1)^2 \right), \text{ and} \quad (19)$$

$$R_\ell = \max \{1, 1 + r_{ss} + \alpha (\pi_\ell - 1)\}. \quad (20)$$

Equations (17) and (18) 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 (18) can be expressed as one equation in the unknown, π_ℓ , after using equations (17) and (19), 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:

$$p = 0.80, r_\ell = -0.0015, G_{ss} = G(r_{ss}) = 0.20, \beta = 0.995,$$

$$\varepsilon = 4, \phi = 110, \chi = 1.25, \alpha = 1.5$$

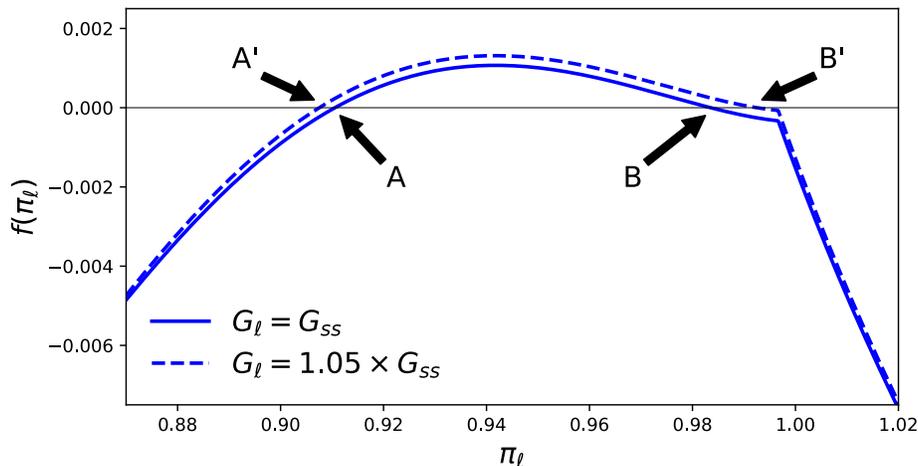
In the $R > 1$ steady-state REE, $C_{ss} = 0.8$, $\pi_{ss} = 1$, $N_{ss} = 1$. While $r = r_\ell$, we set $G_\ell = G(r_\ell) = G_{ss}$. We also consider an alternative specification for government purchases, given by

$$G_\ell = G(r_\ell) = 1.05 \times G(r_{ss}). \quad (21)$$

Figure (1) 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 1 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.

The economy's response to a drop in r is the result of two countervailing forces.

Figure 1: $f(\pi)$ Corresponding to the Target-Inflation Steady-State Equilibrium



Note: The function, f , is defined in the text. The dashed line is discussed in Section 8.1 below. The range of π_ℓ in the figure includes the two values of π_ℓ that correspond to an equilibrium. Source: Authors' calculations.

First, the drop in r leads to an increase in desired savings. In the first best equilibrium, the real interest rate would drop enough to undo the increased desire to save completely, allowing market clearing in the bond and goods market without any change in consumption and employment. When monetary policy is operated by a Taylor rule, and prices are sticky, then we know that policy goes only part-way towards achieving the first best equilibrium. The real interest rate falls, but not by enough so that market clearing must be accomplished in part by a drop in output and income, which reduces the desire to save, as long as the low- r spell is expected to be short enough (that is, p is small enough).¹⁶ If the required fall in the nominal interest rate is sufficiently large, then the ZLB on the nominal interest rate binds. When the ZLB binds, a form of *deflation spiral* is triggered. The fall in output leads to a drop in marginal cost that reduces actual and expected inflation. The latter raises the real interest rate, amplifying the desire to save, leading to an additional drop in actual and expected inflation. An important countervailing force limits the extent of this spiral. As output drops, consumption smoothing leads people to save less. The lower is p , the shorter is the expected duration of the ZLB and the stronger is the consumption smoothing motive.

Three observations about the ZLB follow. First, the logic of the deflation spiral provides intuition into why the fall in output can be very large when the ZLB is binding. The larger the expected deflation in an REE, the larger the drop in output. Second,

¹⁶Further discussion of this point appears below.

Table 1: Equilibrium Values While $r_t = r_\ell$, Returning to Target-Inflation Steady State

Label	Bad ZLB		Good ZLB		Label	Bad ZLB		Good ZLB	
	A	B	A'	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 $G = G_{ss}$ (2a) and when $G = 1.05G_{ss}$ (2b). 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.

the interplay between the deflationary spiral and consumption smoothing provides intuition for why there can be multiple REEs in the ZLB. Third, if p is sufficiently large, the consumption smoothing motive is very weak. When the deflationary spiral is too dominant, an REE does not exist.¹⁷

Turning 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 1). Table 1 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' .

In sum, this section highlights the central role that expected deflation plays in determining the properties of an REE in the ZLB. We expect that because expectations are backward looking, the properties of the learning equilibrium will be very different from those of the REE.

6 Equilibrium Selection

This section considers whether the multiplicity of REEs can be resolved by learnability. We analyze the learnability of an REE by considering a small perturbation in the REE beliefs. We consider these perturbations by analyzing learning equilibria with initial values of Θ that are not REE beliefs but are close to the REE values. We say that an

¹⁷See Werning (2012), who also discusses the possibility of nonexistence of equilibrium in the ZLB.

REE is *learnable* if learning equilibria that begin with beliefs in a neighborhood of the REE beliefs converge to the REE. In this section, we conduct the analysis numerically and consider initial values for Θ that have m_C and m_π equal to $\log(C_\ell)$ and $\log(\pi_\ell)$, respectively, where C_ℓ and π_ℓ are the REE values of C and π . Importantly, the variance of the priors is greater than zero. If learning equilibria starting with these values of Θ converge to the associated REE, then we say the REE is learnable. We have also considered learning equilibria that begin with a vector Θ in which m_C and m_π are near, but not equal to, the associated REE values. In these cases, we find similar results, and our conclusions about learnability remain unchanged.

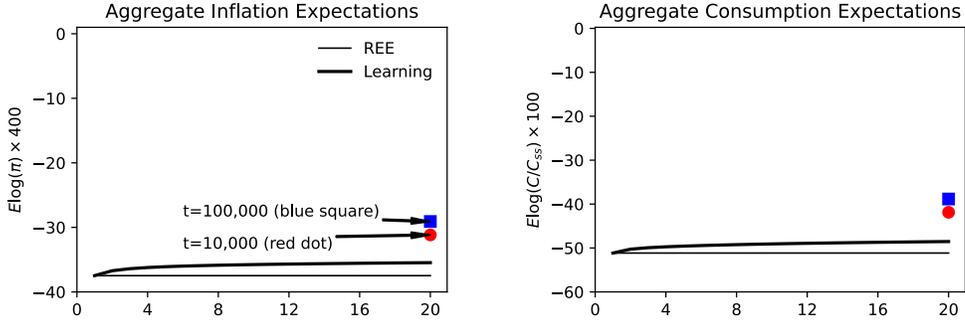
Other initial values of Θ are of particular interest. For example, beliefs with m_C and m_π equal to $\log(C_{ss})$ and $\log(\pi_{ss})$, respectively, are natural candidates in the initial values of Θ . If an REE is learnable and learning equilibria beginning with these initial values of Θ also converge to that REE, then we say that the REE is quasi-globally learnable. Any particular REE cannot be globally learnable in a model with multiple REEs (like the NK model). This result holds because if beliefs are consistent with another REE, then beliefs will not diverge from that equilibrium.

We initially consider the learnability of the Bad-ZLB equilibrium by examining a learning equilibrium with m_i set to the Bad-ZLB equilibrium values. Figure 2a suggests that the learning equilibrium deviates from the Bad-ZLB equilibrium. The red dot shows where that equilibrium is after 10,000 periods and indicates that it is headed toward the Good-ZLB equilibrium. In Section 9, we use linearization methods to prove that at the assumed parameter values, the learning equilibrium cannot converge to the Bad-ZLB equilibrium.¹⁸ We conclude that the Bad-ZLB equilibrium is not learnable.

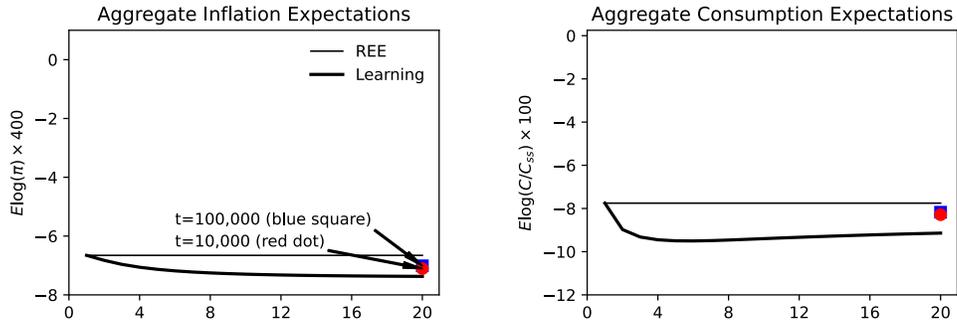
¹⁸Our proof is by contradiction. We linearize our learning model around the Bad-ZLB equilibrium. Suppose the Bad-ZLB equilibrium is stable. Then, the learning equilibrium would eventually (as long as $r = r^\ell$) arrive in an arbitrary small interval, U , about the Bad-ZLB equilibrium, where our linearized system is arbitrarily accurate. We show that that model satisfies the conditions of Theorem 7.2 in Evans and Honkapohja (2001) for beliefs to leave U . This outcome contradicts the hypothesis that the Bad-ZLB equilibrium is stable.

Figure 2: Equilibrium Selection in the ZLB, by Learning

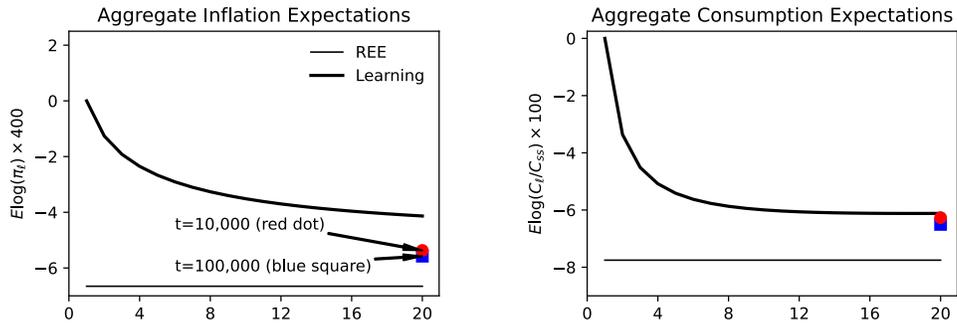
(a) Non-learnability of Bad-ZLB Equilibrium



(b) Learnability of Good-ZLB Equilibrium



(c) Learnability of Good ZLB Equilibrium



Note: In the panels (a) and (b), m_i is initially set to the associated REE value. In panel (c) m_i is initially set to the steady state REE value. In all sub-figures, $\psi_i = 0.02$, $\lambda_i = 1$, $\alpha_i = 2$. Source: Authors' calculations.

We next consider the learnability of the Good-ZLB equilibrium by examining a learning equilibrium with m_i set to the Good-ZLB equilibrium values. Figure 2b shows that the learning equilibrium is converging to the Good-ZLB equilibrium. In Section 9, we use linearization methods to prove that at the assumed parameter values, the

learning equilibrium will converge to the Good-ZLB equilibrium if beliefs start in a neighborhood of that REE. Figure 2c shows that the learning equilibrium converges to the Good-ZLB equilibrium when the beliefs are initially centered on the steady-state REE. These results show that the Good-ZLB is quasi-globally learnable.

7 Speed of Convergence

This section analyzes how show quickly the learning equilibrium converges to the unique learnable REE. In the first subsection, we consider our results for the baseline parameterization of the model. In the second subsection, we consider the effect of the ZLB on the interest rate on the speed of convergence.

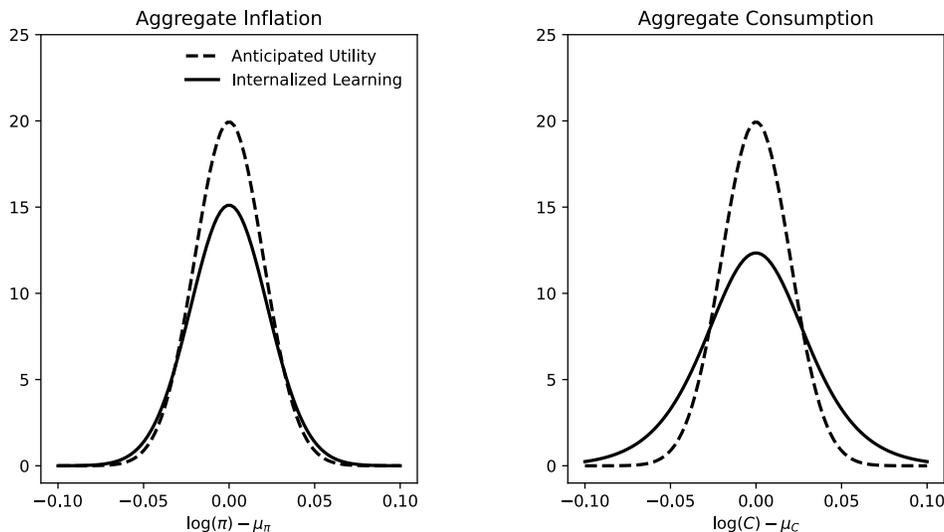
7.1 Baseline Results

Our baseline assumption is that when people are confronted with an unprecedented observation, they become very uncertain about how market-determined variables will evolve. We set the initial value of Θ to the following vector:

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

Figure 3 displays the marginal density of $\log C$ and $\log \pi$ associated with anticipated utility (that is, the Normal distribution evaluated at the prior estimates of the means and variances) and with internalized learning (that is, the marginal data density associated with the truncated Normal-inverse-gamma prior on the parameters of the Normal distribution). Note the fatter tails on the density function associated with internalized learning. The tails are fatter for consumption than inflation because we set a higher upper bound on σ_C (0.05) than on σ_π (0.025). The bounds on the standard deviations correspond to typical period-by-period shock sizes equal to about 6 percent for aggregate consumption and about 10 percentage points for *annualized* aggregate inflation.

Figure 3: Data Density Under Two Models of the Interaction of Beliefs and Decisions



Note: The dashed line corresponds to Normal density functions with means m_i and standard deviations ψ_i . The solid line corresponds to the marginal data density of $\log(x)$ at time one, using Θ before it is updated by the time one value of x is realized. Source: Authors' calculations.

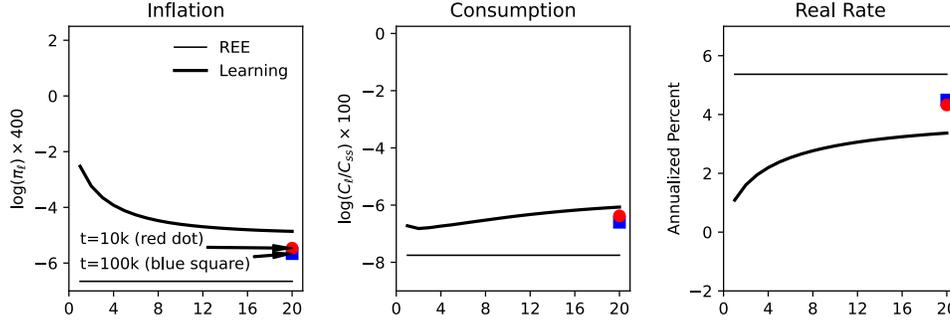
The thin and thick solid lines in Figure 4 display the evolution of inflation, consumption, and the real interest rate after the drop in r under REE and learning, respectively. Consider Figure 4a, which reports results for the REE and internalized learning. Two key features are worth noting. First, in the REE, there is a very large drop in inflation and consumption, and the real interest rate rises sharply. The fall in inflation and consumption and the rise in the real rate are much smaller under learning. Second, the learning economy converges *very* slowly to the REE. After people initially change their views somewhat quickly, the rate at which they change their views slows dramatically (see Figure 2c). For example, the dot labeled $T = 10,000$ displays people's views about the variables after 10,000 quarters. Given our value, $p = 0.8$, r is only expected to be low for about five quarters. Whether convergence to the REE happens after 20 quarters or 10,000 quarters is irrelevant because the ZLB is likely to be over well before that time. The crucial point is that in a typical ZLB episode, people's beliefs are very far from rational expectations, and the associated REE is very different from the learning equilibrium.

Now consider Figure 4b. This figure compares the evolution of the learning equilibrium under anticipated utility (dashed line) and internalized learning (solid line). The key takeaway is that we obtain the same slow-learning result qualitatively regardless of

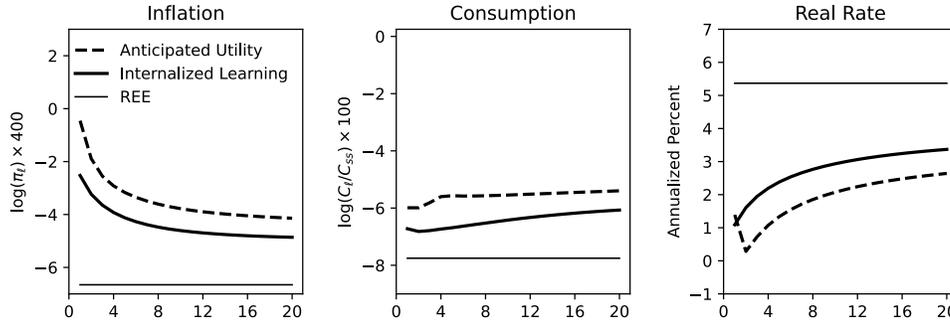
which approach we take to learning. However, consumption and inflation fall somewhat more under internalized learning.

Figure 4: Simulations of Benchmark Model

(a) Speed of Convergence in the Benchmark Model



(b) Comparison, Speed of Convergence Under Anticipated Utility and Benchmark

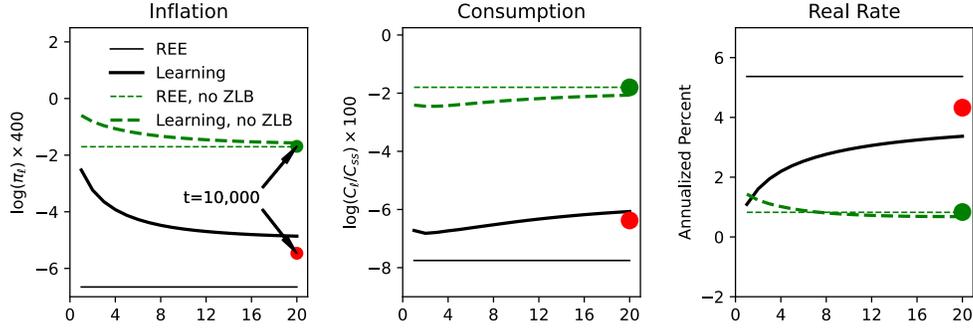


Source: Authors' calculations.

7.2 The Role of the ZLB in the Baseline Results

Figure 5 reports a simulation of our benchmark model in which the ZLB on the interest rate is ignored. For convenience, we reproduce the results from Figure 4 in which the ZLB is binding. The key result is that the learning economy converges very quickly when the ZLB is not binding. The reason is that the Taylor rule weakens the connection between expected and realized inflation. To understand why, suppose people's prior is that inflation will be high in the next period, causing firms to want to raise prices in the current period. When the Taylor principle is operative, the central bank takes action in the current period to lower actual inflation. Because expectations are less self-fulfilling, people will quickly adjust their beliefs. The speed with which they do so depends very much on the value of α , a point that we return to in Section 9.

Figure 5: Benchmark Simulations with and without Binding ZLB



Source: Authors' calculations.

8 Learning and Government Policy

In this section, we analyze the sensitivity of monetary and fiscal policy analysis in the ZLB to deviations from rational expectations. We juxtapose that sensitivity to the lack of sensitivity when the ZLB is not binding.

8.1 Government Purchases Multiplier

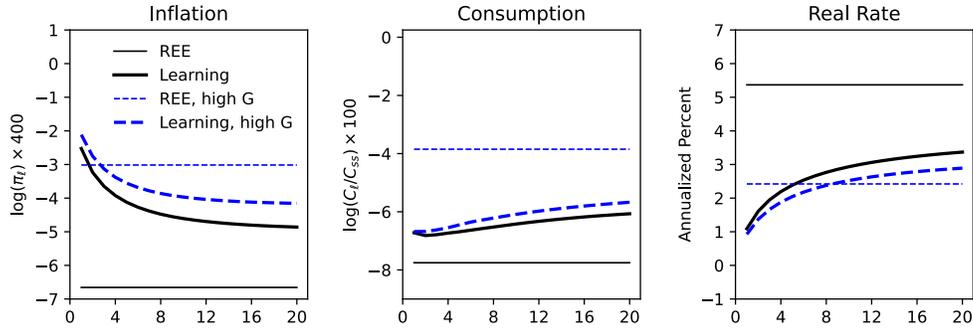
We begin by analyzing the effect of learning on the government purchases multiplier when the ZLB binds. We compute the multiplier by considering the effect on GDP, $C + G$, of a 5 percent rise in government purchases relative to its steady-state level, $G(r_\ell) = 1.05 \times G(r_{ss})$. We denote the difference in consumption and government purchases across the two equilibria by ΔC and $\Delta G = 0.05 \times G(r_{ss})$. Because the Bad-ZLB equilibrium is not stable under learning, we focus on ΔC across Good-ZLB equilibria. We define the multiplier as

$$\frac{\Delta C + \Delta G}{\Delta G}. \quad (22)$$

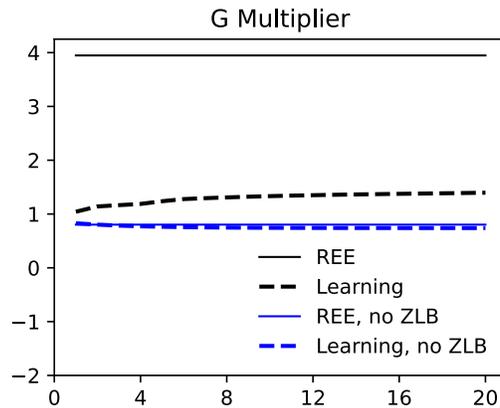
In the REE, the multiplier in the ZLB is equal to 3.95 (see Figure 6b). The multiplier is large when the ZLB is binding because the rise in G generates an increase in expected inflation (see the left panel in Figure 6a). Because R is fixed, this rise generates a fall in the real interest rate and a rise in C (see the middle panel). So, in this case, the multiplier is bigger than one.

Figure 6: Equilibria with and without Jump in G

(a) Increase in Government Purchases During ZLB



(b) Government Purchases Multiplier in ZLB



Notes: The solid lines in Figure 6a reproduce the results based on $G = G_{ss}$ in Figure 4a. The dashed lines report the simulation of the model when $G = 1.05 \times G_{ss}$. Figure 6b displays the government purchases multiplier under internalized learning and in the REE. That figure reports results for the case in which the ZLB is imposed and not imposed ('no ZLB').

Under learning, expected inflation is partially backward-looking and does not move much with a rise in G . As a result, the real interest rate does not fall very much, and the response in consumption is small.

Figure 6b displays the value of the multiplier over time in the REE and under learning in the ZLB. Consistent with the results above, the multiplier under learning is small compared with what it is in the REE. Significantly, the multiplier in the learning equilibrium is very different from the multiplier in the REE over the 20 quarters displayed.

We now turn to the case when the ZLB is not binding. In this case, the REE multiplier, 0.80, is much smaller when the ZLB is binding (see Figure 6b). When the

ZLB is not binding, the rise in inflation causes the monetary authority to raise the real interest rate, which leads to a fall in C . That rise is why the REE multiplier is less than unity outside the ZLB. Figure 6b displays the government purchases multiplier in the learning equilibrium when we ignore the ZLB. Significantly, multiplier values are very similar in the learning equilibrium and the REE. This is consistent with our result that the learning equilibrium converges quite quickly to the REE when the ZLB is not binding.

8.2 Forward Guidance

In this subsection, we consider the sensitivity of the effects of forward guidance to learning. Under such a policy, 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 the first subsection we show that the number of REE proliferates under forward guidance. However, only one of those equilibria is stable under learning. Second, we analyze the effect of forward guidance.

8.2.1 Rational Expectations Equilibria

We construct the REEs with forward guidance by working backward in three steps. First, we compute the unique non-stochastic steady state with $R > 1$. Second, we compute the continuation equilibrium in the period, I , in which r switches from r_ℓ to r_{ss} , where $I \in [2, 3, \dots]$. Third, we compute the equilibrium allocations in the periods before I , denoted by I_{-1} .

In period I , $R = 1$ even though $r = r_{ss}$. People know that the economy will transition to steady state in period $I + 1$. The equilibrium conditions in period I are equations (17) through (19) adjusted for forward guidance:

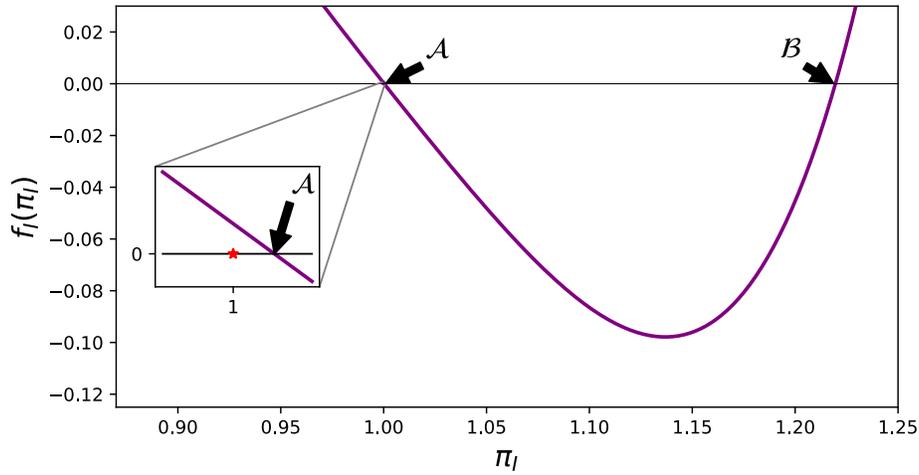
$$1 = \frac{1}{1 + r_{ss}} \frac{C_I}{\pi_{ss} C_{ss}} \quad (23)$$

$$(\pi_I - 1) \pi_I (C_I + G_{ss}) - \frac{\varepsilon - 1}{\phi} (\chi N_I C_I - 1) N_I = 0 \quad (24)$$

$$N_I = (C_I + G_{ss}) \left(1 + \frac{\phi}{2} (\pi_I - 1)^2 \right). \quad (25)$$

Equations (23) and (25) define functions mapping π_I to C_I and N_I . These functions allow us to express the left-hand side of equation (24) 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 7 displays the f_I function for a range of values of π_I . We find two continuation equilibria corresponding to the two zeros of f_I displayed in the figure (see points \mathcal{A} and \mathcal{B}).¹⁹

Figure 7: Equilibria in Period of Switch from $r = r_\ell$ to $r = r_{ss}$ Under One-Period Forward Guidance



Notes: Graph of the function, $f_I(\pi_I)$, discussed after equation (25). The two crossings with the zero line correspond to equilibria in period I , the date when r switches from $r = r_\ell$ to $r = r_{ss}$. 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 equilibrium 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 conditions are the appropriate analog of equations (17) through (19):

$$1 = \frac{1}{1 + r_\ell} \left[p \frac{C_\ell}{\pi_\ell C_\ell} + (1 - p) \frac{C_\ell}{\pi_I C_I} \right] \quad (26)$$

¹⁹From equation (23) 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 7 are the only zeros of f_I .

$$\begin{aligned}
& (\pi_\ell - 1) \pi_\ell (C_\ell + G_\ell) - \frac{\varepsilon - 1}{\phi} (\chi N_\ell C_\ell - 1) N_\ell \\
& - \frac{1}{1 + r_\ell} \left[p (\pi_\ell - 1) \pi_\ell (C_\ell + G^\ell) + (1 - p) (\pi_I - 1) \pi_I \frac{C_\ell}{C_I} (C_I + G_{ss}) \right] = 0 \quad (27)
\end{aligned}$$

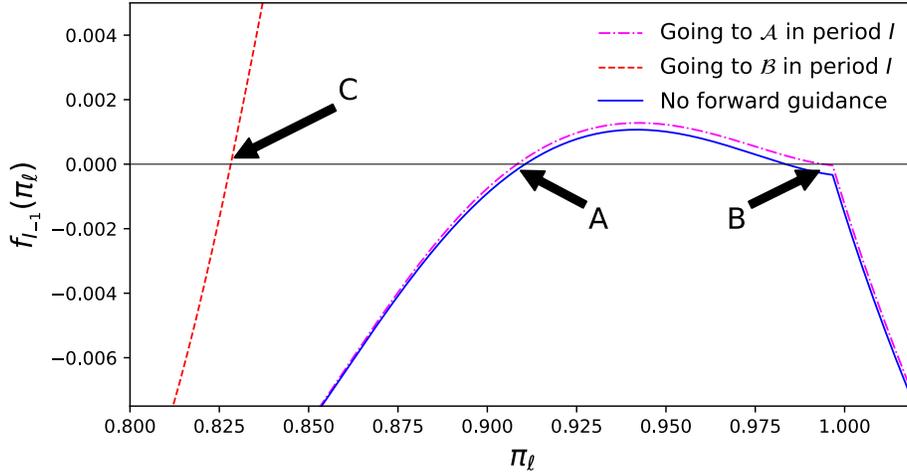
$$N_\ell = (C_\ell + G^\ell) \left(1 + \frac{\phi}{2} (\pi_\ell - 1)^2 \right) \quad (28)$$

Here, we impose the condition that $R_\ell = 1$. In effect, we assume that the ZLB is binding in periods I_{-1} , and the Taylor rule holds. This assumption is satisfied in all of the examples we have studied.

Given C_I and π_I , equations (26) through (28) define a mapping from π_ℓ to C_ℓ and N_ℓ . Now, we can express the left-hand side of equation (27) 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 8 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 8) and one associated with \mathcal{B} (see C in Figure 8). So there are three REEs with forward guidance. The two REEs without forward guidance can be seen in the solid line in Figure 8 (we take this curve from Figure 1).

Figure 8: REE Equilibria at the ZLB with and without Forward Guidance



Notes: The solid line reproduces the solid line in Figure 1 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 to r_{ss} (that is, 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 7).

8.2.2 Learning Equilibria

In the period of forward guidance, $r = r_{ss}$, $R = 1$. In all periods when $r = r_\ell$ (that is, I_{-1}), people understand that the economy reverts to an REE when $r = r_{ss}$. As discussed, there are two REEs starting in period I , the first date when $r = r_{ss}$ (see points \mathcal{A} and \mathcal{B} in Figure 7).

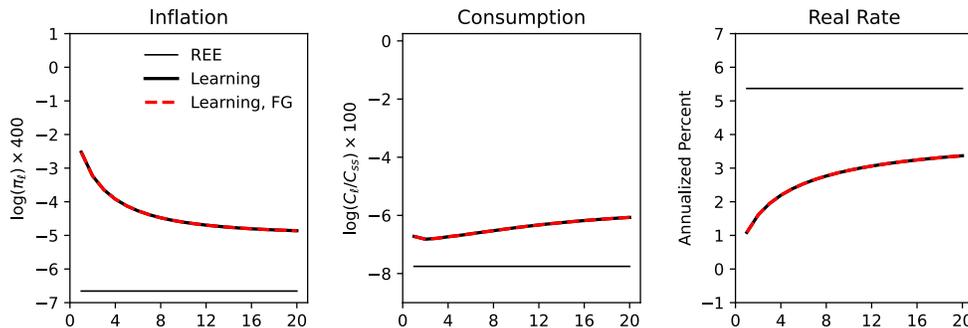
We begin by analyzing whether any of the learning equilibria converge to a particular REE in I_{-1} . Two of the three REEs in I_{-1} are not learnable. These are the equilibria associated with points \mathcal{A} and \mathcal{C} in Figure 8. In contrast, the equilibrium represented by \mathcal{B} is learnable. Thus, learnability selects a unique REE.

We next consider the learning equilibrium using the same initial values for Θ as in Section 7.1. The learning equilibrium under forward guidance is indistinguishable from the learning equilibrium without forward guidance, as shown in Figure 9. The reason is that forward guidance has small effects on inflation and consumption in period I . In turn, these small effects have almost no effect on the learning equilibrium where beliefs about π_ℓ and C_ℓ are partially backward looking. Clearly, there is no forward guidance puzzle under learning.²⁰ That puzzle emerges under RE because of the strong effect of

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

forward guidance on expected inflation. Under learning, expectations are backward-looking, and forward guidance has little influence on expected inflation. Note that the learning equilibria with and without forward guidance converge slowly.

Figure 9: Forward Guidance Under Learning and REE



9 The value of b in the NK Model

In this section, we accomplish two tasks. First, we demonstrate that b (from propositions 1 and 2) is large in a linearized solution of our NK model when the ZLB is binding. Second, we argue that the asymptotic rate of convergence is a good guide to the small t rate of convergence.

We base our analysis below on linearized versions of the policy functions defined using equations (5) and (12). Here, we find using time notation (rather than recursive notation) convenient. The details of our linearization appear in Appendix C. Recall that the household problem can be reduced to finding an optimal decision rule for bond holdings, $b'(b_h, \Theta, x)$, denoted here by $b_{h,t}$. Log-linearizing this decision rule, we obtain

$$\hat{b}_{h,t} = \gamma_{b,b} \hat{b}_{h,t-1} + \gamma_{b,\pi} \hat{\pi}_t + \gamma_{b,C} \hat{C}_t + \gamma_{b,\mu_\pi} \hat{m}_{\pi,t} + \gamma_{b,\mu_C} \hat{m}_{C,t}. \quad (29)$$

With one exception, the hat notation, \hat{q}_t , denotes $\log(q_t/q)$, and q is the REE value of q_t about which the linearization is done. The exception is household bond holdings, $b_{h,t}$, where $\hat{b}_{h,t} = b_{h,t} - b_h$. Also, $\hat{\mu}_t = [\hat{m}_{\pi,t}, \hat{m}_{C,t}]$ represents the log deviation of people's time t posterior of $\mathbb{E}_t x_{t+1}$ and the REE value of Ex_{t+1} conditional on $r_{t+1} =$

meaningful.

r^ℓ .²¹ We use the posterior means ($\mathbb{E}_t x_{t+1}$) rather than the prior means ($\mathbb{E}_{t-1} x_{t+1}$) because households and firms can compute Θ' using Θ and x . Variances of beliefs do not appear because of the certainty equivalence implied by the linearization. The parameters in equation (29) are functions of model parameters and the point about which the linearization is done. The linearized price decision rule, $p'_f(p_f, \Theta, x)$, of the firm (denoted by $\hat{p}_{f,t}$) is given by

$$\hat{p}_{f,t} = \gamma_{p,p} \hat{p}_{f,t-1} + \gamma_{p,\pi} \hat{\pi}_t + \gamma_{p,C} \hat{C}_t + \gamma_{p,\mu_\pi} \hat{m}_{\pi,t} + \gamma_{p,\mu_C} \hat{m}_{C,t}. \quad (30)$$

The time t realized value of \hat{x}_t enters the decision rules, equations (29) and (30), by two channels. The first channel reflects that people use \hat{x}_t to determine the period t values of the exogenous variables in their period t budget constraint. The second channel reflects that $\hat{\mu}_t$ depends on \hat{x}_t , $\hat{\mu}_{t-1}$, and the gain in the Bayesian updating equation.

In each period, we compute a linearized period equilibrium (see Definition 2), so that (i) $\hat{b}_{h,t-1} = \hat{p}_{f,t-1} = 0$ and (ii) \hat{x}_t is determined by the requirements, $\hat{b}_{h,t} = 0$ and $\hat{p}_{f,t} = 0$:

$$\begin{aligned} 0 &= \gamma_{b,\pi} \hat{\pi}_t + \gamma_{b,C} \hat{C}_t + \gamma_{b,\mu_\pi} \hat{m}_{\pi,t} + \gamma_{b,\mu_C} \hat{m}_{C,t} \\ 0 &= \gamma_{p,\pi} \hat{\pi}_t + \gamma_{p,C} \hat{C}_t + \gamma_{p,\mu_\pi} \hat{m}_{\pi,t} + \gamma_{p,\mu_C} \hat{m}_{C,t}. \end{aligned}$$

Assuming the relevant matrix inverse exists, \hat{x}_t is given by

$$\hat{x}_t = B \hat{\mu}_t, \quad (31)$$

where²²

$$B = - \begin{bmatrix} \gamma_{b,\pi} & \gamma_{b,C} \\ \gamma_{p,\pi} & \gamma_{p,C} \end{bmatrix}^{-1} \begin{bmatrix} \gamma_{b,\mu_\pi} & \gamma_{b,\mu_C} \\ \gamma_{p,\mu_\pi} & \gamma_{p,\mu_C} \end{bmatrix}.$$

The law of motion of $\hat{\mu}_t$ is a stacked version of the updating expressions in equation (??). For simplicity, we assume $\gamma_t = 1/(\lambda_0 + t)$ in both equations. Here, λ_0 denotes

²¹Note that the only exogenous random variable in our model is the natural rate of interest. Given the model structure, in the REE, $E(x_{t+1}|r_{t+1} = r^\ell)$ is equal to the value of x_t while $r_t = r^\ell$. We discuss posterior and prior means because when the natural rate of interest is low households and firms believe that there are additional sources of variation (see equation 15). See Mayer (2021) for a discussion of rates of convergence in univariate learning models with additional stochastic regressors.

²²In the examples that we have considered, we have not encountered an exception to the invertibility assumption.

the initial precision of beliefs about the mean of inflation and consumption. Combining the vector Bayesian updating expression with equation (31), we obtain:

$$\hat{\mu}_t = \hat{\mu}_{t-1} + \gamma_t [B(1 - \gamma_t)(I - \gamma_t B)^{-1} \hat{\mu}_{t-1} - \hat{\mu}_{t-1}]. \quad (32)$$

This equation is of the same form as equation 1, with

$$M(\hat{\mu}_{t-1}, \gamma_t) = B(1 - \gamma_t)(I - \gamma_t B)^{-1} \hat{\mu}_{t-1}.$$

So, propositions 1 and 2 apply to this system.

Table 2 displays the eigenvalues of $D_1 M(0, 0)$ corresponding to the Good- and Bad - ZLB equilibria for the benchmark parameter values. The maximal eigenvalue ('Eigenvalue 1'), b , associated with the Good-ZLB and Bad-ZLB equilibria are 0.92 and 1.26, respectively. Consistent with Section 6, the Bad-ZLB equilibrium is not locally learnable because $b > 1$. The Good-ZLB equilibrium is locally learnable because the corresponding value of b is less than one.

Asymptotic convergence to the Good-ZLB REE is slow because b is large. According to Propositions 1 and 2, the asymptotic rate of convergence of $\|\hat{\mu}_t\|$ is t^{b-1} . Using t^{b-1} as an approximation to $\|\hat{\mu}_t\|$, the amount of time, T_1 , it takes to close two-thirds of a gap, $\|\hat{\mu}_0\|$, is given by $T_1 = (1/3)^{\frac{1}{b-1}}$. Table 2 reports values of T_1 for different variants of the model. In the benchmark model where the ZLB is binding, $b = 0.92$ and $T_1 = 920,482$. This large value of T_1 is qualitatively consistent with the basic prediction of the non-linear solution to the model—namely that the rate of convergence is quite slow (see Figure 4a). Similarly, the small value of T_1 reported in the table for the case in which the ZLB is not binding and $\alpha = 1.5$ is qualitatively consistent with the finding for the non-linear solution to the model (see Figure 5). In this sense, the asymptotic results in Propositions 1 and 2 are useful guides for the rate of convergence when t is small.

A different way to assess the usefulness of the asymptotics is to calculate the actual amount of time, T , required to close two-thirds of the initial gap between priors and steady state according to the linearized solution to the model.²³ To this end, we simulate the linearized solution to the model when the ZLB is binding and when we

²³The initial gap in $\log x_i$, $i = 1, 2$, corresponds to the log-deviation of x_i in the initial steady state and the REE while $r = r_\ell$.

ignore the ZLB. In the latter case, we consider $\alpha = 1.5$ and 3 . The results are reported in Table 2. We find that, for the benchmark model, when the ZLB is binding, $T = 944,710$. In sharp contrast, when the ZLB is not binding and $\alpha = 1.5$ and 3 , we find $T = 3$ and 1 periods, respectively. The importance of the ZLB and the value of α in determining the speed of convergence are qualitatively the same as our results using T_1 .

Table 2: Eigenvalues of B

	Eigenvalue 1	Eigenvalue 2	T_1	T
Good ZLB	0.92	-0.48	920,482	944,710
Bad ZLB	1.26	-1.21	NA	NA
No ZLB, $\alpha = 1.5$	0.054+0.44i	0.054-0.44i	2	3
No ZLB, $\alpha = 3$	-0.135+0.84i	-0.135-0.84i	2	1

Note: The matrix, B , is defined in equation (31). The scalar, b , discussed in the text is the largest real part of the eigenvalues of B . The reported values of T are based on simulations of the linearized solution to the model. For the definitions of T and T_1 see the text.

10 Conclusion

In this paper, we consider the speed with which people learn about their environment after an unusual event. We do so in a non-linear NK model with internally rational households and firms that are learning about how the economy will evolve after the event. To characterize the speed of convergence of people's beliefs, we analytically extend results in the literature to characterize the asymptotic convergence rate of multivariate systems. We argue that the slow convergence result arises naturally in the NK model when the ZLB is binding. Under these circumstances, analyses of fiscal and monetary policies under rational expectations can be very misleading. It would be interesting to pursue this possibility in an empirically plausible version of the NK model of the sort considered by Christiano et al. (2016) or Del Negro et al. (2023).

Finally, we note that there are other circumstances in which slow learning could arise. For example, plausible parameterizations of Cagan (1956)'s model of money demand under hyperinflation map into high b economies. Results in Marcet and Sargent (1995, Table 6.3) imply that estimates of the elasticity of money demand in hyperinflations (see, for example, Christiano (1987) and Taylor (1991)) map into high values of b . More generally, our results suggests that any model with strong strategic complementarities could exhibit slow convergence to rational expectations under 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 Chen Lian**, “Dampening general equilibrium: From micro to macro,” Technical Report, National Bureau of Economic Research 2017.
- **and** – , “Forward Guidance without Common Knowledge,” *American Economic Review*, September 2018, 108 (9), 2477–2512.
- Arifovic, Jasmina, Stephanie Schmitt-Grohé, and Martin Uribe**, “Learning to live in a liquidity trap,” *Journal of Economic Dynamics and Control*, 2018, 89, 120–136.
- Benhabib, Jess, Stephanie Schmitt-Grohé, and Martin Uribe**, “Monetary policy and multiple equilibria,” *American Economic Review*, 2001, 74 (2), 165–170.
- Cagan, Phillip**, “The monetary dynamics of hyperinflation,” *Studies in the Quantity Theory of Money*, 1956.
- Christiano, Lawrence J**, “Cagan’s model of hyperinflation under rational expectations,” *International Economic Review*, 1987, pp. 33–49.
- Christiano, Lawrence J., Martin S. Eichenbaum, and Mathias Trabandt**, “Unemployment and Business Cycles,” *Econometrica*, 2016, 84 (4), 1523–1569.
- , – , **and Sergio Rebelo**, “When is the Government Spending Multiplier Large?,” *Journal of Political Economy*, 2011, (February).
- Christopeit, Norbert and Michael Massmann**, “Estimating structural parameters in regression models with adaptive learning,” *Econometric Theory*, 2018, 34 (1), 68–111.
- Cogley, Timothy and Thomas J Sargent**, “Anticipated utility and rational expectations as approximations of Bayesian decision making,” *International Economic Review*, 2008, 49 (1), 185–221.
- Eggertsson, Gauti and Michael Woodford**, “The Zero Bound on Interest Rates and Optimal Monetary Policy,” *Brookings Papers on Economic Activity*, 2003, (1).
- Eggertsson, Gauti B and Michael Woodford**, “Policy options in a liquidity trap,” *American Economic Review*, 2004, 94 (2), 76–79.

- Erceg, Christopher J. and Andrew T. Levin**, “Imperfect credibility and inflation persistence,” *Journal of Monetary Economics*, 2003, 50 (4), 915–944.
- Eusepi, Stefano, Christopher Gibbs, and Bruce Preston**, “Monetary policy trade-offs at the zero lower bound,” *Available at SSRN 4069078*, 2022.
- Evans, George W**, “Theories of learning and economic policy,” *Revue d’économie politique*, 2021, 132 (3-4), 583–608.
- **and Seppo Honkapohja**, “Local convergence of recursive learning to steady states and cycles in stochastic nonlinear models,” *Econometrica*, 1995, pp. 195–206.
- **and –**, “Convergence for difference equations with vanishing time-dependence, with applications to adaptive learning,” *Economic Theory*, 2000, 15, 717–725.
- Evans, George W. and Seppo Honkapohja**, *Learning and Expectations in Macroeconomics*, Princeton University Press, 2001.
- Farhi, E, M Petri, and I Werning**, “The fiscal multiplier puzzle: Liquidity traps, bounded rationality, and incomplete markets,” Technical Report, mimeo 2020.
- Farhi, Emmanuel and Iván Werning**, “Monetary policy, bounded rationality, and incomplete markets,” *American Economic Review*, 2019, 109 (11), 3887–3928.
- Farmer, Leland, Emi Nakamura, and Jón Steinsson**, “Learning about the long run,” Technical Report, National Bureau of Economic Research 2021.
- Ferrero, Giuseppe**, “Monetary policy, learning and the speed of convergence,” *Journal of Economic Dynamics and Control*, 2007, 31 (9), 3006–3041.
- García-Schmidt, Mariana and Michael Woodford**, “Are low interest rates deflationary? A paradox of perfect-foresight analysis,” *American Economic Review*, 2019, 109 (1), 86–120.
- Gust, Christopher, Edward Herbst, and David López-Salido**, “Forward guidance with bayesian learning and estimation,” 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.
- Hommes, Cars**, “The heterogeneous expectations hypothesis: Some evidence from the lab,” *Journal of Economic dynamics and control*, 2011, 35 (1), 1–24.
- Iovino, Luigi and Dmitriy Sergeyev**, “Central bank balance sheet policies without rational expectations,” *Review of Economic Studies*, 2023, 90 (6), 3119–3152.

- Kaplan, Greg and Giovanni L Violante**, “Microeconomic heterogeneity and macroeconomic shocks,” *Journal of Economic Perspectives*, 2018, 32 (3), 167–94.
- Kreps, David M**, “Anticipated utility and dynamic choice,” *Econometric Society Monographs*, 1998, 29, 242–274.
- Ljung, Lennart**, “Analysis of recursive stochastic algorithms,” *IEEE transactions on automatic control*, 1977, 22 (4), 551–575.
- Marcet, Albert and Thomas J Sargent**, “Convergence of least-squares learning in environments with hidden state variables and private information,” *Journal of Political Economy*, 1989, 97 (6), 1306–1322.
- **and** —, “Convergence of least squares learning mechanisms in self-referential linear stochastic models,” *Journal of Economic theory*, 1989, 48 (2), 337–368.
- **and Thomas Sargent**, “Speed of Convergence of Recursive Least Squares: Learning with Autoregressive Moving-Average Perceptions,” in Alan Kirman and Mark Salmon, eds., *Learning and rationality in economics*, Oxford University Press, UK, 1995, chapter 6, pp. 179–215.
- Negro, Marco Del, Marc P Giannoni, and Christina Patterson**, “The forward guidance puzzle,” *Journal of Political Economy Macroeconomics*, 2023, 1 (1), 43–79.
- Preston, Bruce**, “Adaptive learning in infinite horizon decision problems,” *Unpublished, Columbia University*, 2005.
- Rotemberg, Julio J**, “Sticky prices in the United States,” *Journal of political economy*, 1982, 90 (6), 1187–1211.
- Taylor, Mark P**, “The hyperinflation model of money demand revisited,” *Journal of money, Credit and Banking*, 1991, 23 (3), 327–351.
- Vimercati, Riccardo Bianchi, Martin S Eichenbaum, and Joao Guerreiro**, “Fiscal policy at the zero lower bound without rational expectations,” Technical Report, National Bureau of Economic Research 2021.
- Vives, Xavier**, “How fast do rational agents learn?,” *The Review of Economic Studies*, 1993, 60 (2), 329–347.
- Werning, Ivan**, “Managing a Liquidity Trap: Monetary and Fiscal Policy, manuscript,” 2012.
- Woodford, Michael**, “Self-fulfilling expectations and fluctuations in aggregate demand,” 1990.

- , “Methods of Policy Accommodation at the Interest-rate Lower Bound,” *Proceedings-Economic Policy Symposium-Jackson Hole, Federal Reserve Bank of Kansas City*, 2012, pp. 185–288.
- **and Yinxu Xie**, “Policy options at the zero lower bound when foresight is limited,” in “AEA Papers and Proceedings,” Vol. 109 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 2019, pp. 433–437.
- **and** – , “Fiscal and monetary stabilization policy at the zero lower bound: Consequences of limited foresight,” *Journal of Monetary Economics*, 2022, *125*, 18–35.