

NBER WORKING PAPER SERIES

NEURAL NETWORK LEARNING FOR NONLINEAR ECONOMIES

Julian Ashwin
Paul Beaudry
Martin Ellison

Working Paper 32807
<http://www.nber.org/papers/w32807>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2024

The authors gratefully acknowledge discussions with Guido Ascari, Roger Farmer, Friedrich Geiecke, Thomas Lubik, Michael McMahon, Leonardo Melosi, Galo Nuño, Ernesto Pastén, Michael Reiter, and useful comments from participants at the International Symposium in Computational Economics and Finance, the Young Economist Symposium (YES), the Warsaw International Economics Meeting, the Econometric Society Winter Meeting, the Society for Nonlinear Dynamics and Econometrics Symposium and the SNDE session at the ASSA meetings. We also thank an anonymous referee for making numerous suggestions that improved the paper. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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/w32807>

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.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 Julian Ashwin, Paul Beaudry, and Martin Ellison. 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.

Neural Network Learning for Nonlinear Economies
Julian Ashwin, Paul Beaudry, and Martin Ellison
NBER Working Paper No. 32807
August 2024
JEL No. C45,E19,E47

ABSTRACT

Neural networks offer a promising tool for the analysis of nonlinear economies. In this paper, we derive conditions for the global stability of nonlinear rational expectations equilibria under neural network learning. We demonstrate the applicability of the conditions in analytical and numerical examples where the nonlinearity is caused by monetary policy targeting a range, rather than a specific value, of inflation. If shock persistence is high or there is inertia in the structure of the economy, then the only rational expectations equilibria that are learnable may involve inflation spending long periods outside its target range. Neural network learning is also useful for solving and selecting between multiple equilibria and steady states in other settings, such as when there is a zero lower bound on the nominal interest rate.

Julian Ashwin
School of Business and Economics
Maastricht University
Tongersestraat 53
Maastricht, 6211 LM
Netherlands
julian.ashwin@maastrichtuniversity.nl

Martin Ellison
Department of Economics
University of Oxford
Manor Road Building
Oxford
OX1 2UQ
United Kingdom
martin.ellison@economics.ox.ac.uk

Paul Beaudry
Vancouver School of Economics,
University of British Columbia
6000 Iona Drive
Vancouver, V6T 1Z1
Canada
and NBER
Paul.Beaudry@ubc.ca

1 Introduction

Nonlinearities abound in macroeconomics. They arise whenever the structure of the economy is nonlinear, as in the flattening of the Phillips curve at low inflation (Forbes et al., 2021; Benigno and Eggertsson, 2023), or if policy itself is nonlinear. A well-known example of the latter is the zero lower bound on nominal interest rates, which has the potential to generate multiple steady states and regions of local indeterminacy (Benhabib et al., 2001; Eggertsson and Woodford, 2003). The richness of the dynamics in these settings is challenging for traditional analysis based on perturbations of perfect foresight equilibria, and has stimulated interest in using techniques from machine learning to understand better the stability and dynamics of rational expectations equilibria. Neural network learning is at the forefront of such efforts (Maliar et al., 2021; Kase et al., 2022).

This paper breaks new ground in the analysis of nonlinear economies. It is widely understood that a sufficiently flexible neural network can approximate any real-valued continuous function arbitrarily well (Cybenko, 1989; Hornik et al., 1989; Hornik, 1991), so any nonlinear Rational Expectations Equilibrium (REE) can in principle be represented by a neural network.¹ We ask whether agents can learn this neural network representation. In linear models, this is already covered by the T-map and e-stability conditions popularised by Evans and Honkapohja (2001). For nonlinear models, we show how stochastic approximation theory can be used to derive conditions for the stability of a nonlinear REE under neural network learning. Our results highlight the potential of neural network learning as a tool for the analysis of nonlinear economies.

We show the value of neural network learning when monetary policy targets a range, rather than a specific value, of inflation. At its meeting in September 2019, the FOMC discussed the adoption of such a target range for inflation when deliberating on the future of the US monetary policy framework.² Further back, Governor Ben S. Bernanke remarked in 2003 that publication of a target range for inflation was the “hallmark of inflation targeting.”³ Recent research by Bianchi et al. (2021) and Le Bihan et al. (2023) has begun exploring some of the implications of introducing a target range for inflation. The implicit assumption throughout is that a target range induces nonlinearities because monetary policy reacts differently depending on whether inflation is below, above, or within its target range.

Our results highlight the importance of the persistence of shocks and inertia in the structural equations of the economy. If these are low then a unique REE exists with inflation spending prolonged periods in its target range and a neural network representation that agents can learn. However, if persistence or inertia is high then multiple REE exist, only some of which have a neural network representation that agents can learn. In particular, the neural network representation of REEs with inflation spending long periods in the target range is not learnable whereas those with inflation mostly outside its target range are. The implication is that inflation will rarely be in the target range if the shock persistence or structural inertia is high.

The paper improves on the existing literature by demonstrating how neural network learning can inform

¹A neural network representation exists for any determinate or indeterminate REE.

²Minutes at <https://www.federalreserve.gov/monetarypolicy/files/fomcminutes20190918.pdf>.

³“Constrained Discretion” and Monetary Policy, <https://www.federalreserve.gov/BOARDDOCS/Speeches/2003/20030203/default.htm>.

the global stability analysis of REE in stochastic nonlinear environments. The closest antecedents to our work are Evans et al. (2008) and Evans et al. (2022), who show that adaptive learning beliefs formed in a sequence of non-rational temporary equilibria diverge from an unstable steady state. Expectations in these equilibria will never be consistent with rationality, as in other papers that deal with nonlinearity by modelling linear learning in a nonlinear model (Hommes and Sorger, 1998; Bullard, 1994). Our application of neural network learning is instead the first to permit global analysis of nonlinear models with fully rational expectations.

We study neural network learning with a target range for inflation or a zero lower bound on the nominal interest rate, but there are many other examples of stochastic frameworks that also have multiple steady states, some of which are locally indeterminate. Our analysis speaks to all these environments, in particular questioning the validity of conclusions based on dynamics being close to the indeterminate steady state for long periods. Noteworthy examples include increasing returns to scale in a production function (Cazzavillan et al., 1998); externalities in a search and matching framework for goods (Kaplan and Menzio, 2016) or labour (Eeckhout and Lindenlaub, 2019); endogenous markups generating non-convex marginal revenue product of capital (Galí, 1995); complementarities in R&D investment (Greiner and Bondarev, 2017); correlated private information about asset payoffs (Manzano and Vives, 2011); non-convex relationships between economic activity and ecological systems (Mäler et al., 2003); positive externalities in production (Krugman, 1991); feedback between government debt and interest rates through liquidity constraints (Angeletos et al., 2021).

The structure of the paper is as follows: Section 2 introduces neural network learning and shows how stochastic approximation theory can be used to assess the stability of a nonlinear REE under learning. An application where monetary policy has a target range for inflation is in Section 3. The environment is nonlinear but purely forward-looking, so the stability of REE under neural network learning can be characterised analytically even though shocks are persistent. Section 4 shows that much of the intuition from the analytic results carries over to a numerical example with backward-looking inertia in the structure of the economy, and Section 5 warns against overinterpreting empirical tests for indeterminacy. The analysis is extended to an environment with a zero lower bound on the nominal interest rate in Section 6. A final section concludes.

2 Neural networks

Neural networks offer a promising tool for the analysis of learning in nonlinear economies because they allow agents to approximate arbitrarily complex nonlinear functions in a way that is efficient to estimate. This Section begins by describing a generic single-layer feedforward neural network. We then assume that agents use the neural network recursively to learn and form expectations in the nonlinear economy. Our insight recognises that such neural network learning describes a stochastic recursive algorithm, allowing us to derive the conditions for agents to learn a nonlinear rational expectations equilibrium.

2.1 A single-layer feedforward neural network

The neural network in Figure 1 uses a $1 \times n$ vector of inputs X to predict a scalar output y . At the input layer, linear combinations of the X 's are created using the parameters in w_1 to form H . If there are n inputs

and K nodes in the neural network then w_1 is an $n \times K$ matrix and H is a $1 \times K$ vector. In the hidden layer, the linear combinations H are fed through activation functions $\psi(H)$, which at their simplest are rectified linear activation functions (ReLU). For each element H_i in the vector H , these functions return the value H_i if $H_i > 0$, otherwise they return zero. We say that a node is activated when $H_i > 0$ so that input H_i is passed to A_i . At the output layer, a linear combination of the $1 \times K$ vector A from the activation functions is created using a $K \times 1$ vector of parameters w_2 to form a prediction $\hat{y}$ of scalar output y .




Figure 1: A single-layer feedforward neural network

The network in Figure 1 is single-layered and feedforward because there is only one hidden layer and there are no feedback connections between the layers. In the neural network literature, the parameters in w_1 and w_2 are referred to as “biases” if they relate to a constant term in X and “weights” otherwise. According to the Universal Approximation Theorem of Cybenko (1989), a neural network of this type can approximate any real-valued continuous function arbitrarily well, provided there are sufficient nodes in the hidden layer. When the activation function is ReLU, the theorem can be understood as the neural network fitting a piecewise linear function with enough linear pieces. The slope of each linear piece depends on the slope of the combinations of inputs that pass through the activation functions, and the breakpoints between linear pieces are identified where the activation functions turn on and off. A single-layered feedforward neural network thus has the potential to capture the expectations formation process in any rational expectations equilibrium.

2.2 Neural network learning

We assume that agents use a neural network to form expectations in the nonlinear economy. They learn the parameters of the network by the method of gradient descent, which updates their estimates in response to their latest prediction error. The direction of each update is determined by the gradient (derivative) of the prediction error with respect to the parameters of interest, ensuring that each parameter is updated in a direction that improves the accuracy of predictions. The magnitude of each update is regulated by a learning rate parameter, $\gamma_t = 1/t$.

The method of gradient descent updates parameter estimates according to:

$$w'_1 = w_1 - \gamma_t \nabla_{w_1} \zeta_t, \quad (1)$$

$$w'_2 = w_2 - \gamma_t \nabla_{w_2} \zeta_t, \quad (2)$$

where ∇_{w_1} and ∇_{w_2} are partial derivatives of the prediction error $\zeta_t = y_t - \psi(X_t w_1) w_2$.

2.3 Stochastic approximation

Equations (1) and (2) describe a stochastic recursive algorithm of the standard form:

$$\underbrace{\begin{bmatrix} w'_1 \\ w'_2 \end{bmatrix}}_{\theta_t} = \underbrace{\begin{bmatrix} w_1 \\ w_2 \end{bmatrix}}_{\theta_{t-1}} + \gamma_t \underbrace{\begin{bmatrix} X_t^T w_2^T \psi'(X_t w_1) \\ \psi(X_t w_1)^T \end{bmatrix}}_{Q(\theta_{t-1}, Z_t)} \zeta_t, \quad (3)$$

where θ_t is a vector of parameter estimates, Z_t is a vector of inputs X_t and output y_t , and $Q(\theta_{t-1}, Z_t)$ is a matrix of updates obtained from the prediction error and its partial derivatives.

Stochastic approximation theory studies the behaviour of recursive algorithms such as ours. An accessible introduction for macroeconomists is in Section 2.7 of Evans and Honkapohja (2001), where they discuss how a standard recursive stochastic algorithm can be associated with an ordinary differential equation that shares many of its characteristics. In particular, the recursive algorithm and differential equation have the same limit points and stability properties, which allows us to derive analytical conditions for the existence and stability of a rational expectations equilibrium under neural network learning.

The associated differential equation describes the updating of parameter estimates θ in notional time τ :

$$\frac{d\theta}{d\tau} = h(\theta(\tau)), \quad (4)$$

$$h(\theta) = \lim_{t \rightarrow \infty} EQ(\theta, \bar{Z}_t(\theta)), \quad (5)$$

where $h(\theta(\tau))$ describes the expected update as a function of $\bar{Z}_t(\theta)$, the stochastic process for Z_t obtained when holding θ_{t-1} at the fixed value θ . The differential equation has a limit point when $h(\bar{\theta}) = 0$ and is asymptotically stable if all the eigenvalues of the Jacobian $Dh(\bar{\theta})$ have negative real parts. Such a limit point represents an equilibrium of the system; it is learnable if it is asymptotically stable.

3 Analytic example with target range for inflation

This Section analyses neural network learning in an example nonlinear economy. The setting is a New Keynesian model in which monetary policy only reacts when inflation is outside a target range. The source of the nonlinearity is the policy rule for the nominal interest rate, which behaves differently when inflation is below, in, or above the target range. We present the model and its rational expectations equilibrium, and calculate what happens if agents form expectations using neural network learning. Derivations are in Appendix A.

3.1 Environment

Inflation π_t and output y_t are determined by a New Keynesian Phillips curve, a static IS curve⁴, and a policy rule for the nominal interest rate R_t that reacts to inflation with coefficient $\alpha > 1$ when inflation is outside the target range $[-\pi^*, \pi^*]$.

$$\pi_t = E_t \pi_{t+1} + \kappa y_t + \epsilon_t \quad (6)$$

$$y_t = -\sigma(R_t - E_t \pi_{t+1}) \quad (7)$$

$$R_t = \begin{cases} \alpha(\pi_t + \pi^*) & \text{if } \pi_t < -\pi^* \\ 0 & \text{if } -\pi^* \leq \pi_t \leq \pi^* \\ \alpha(\pi_t - \pi^*) & \text{if } \pi_t > \pi^* \end{cases} \quad (8)$$

The disturbance term ϵ_t is assumed to follow an n -state exogenous Markov chain, with transitions between states governed by an $n \times n$ probability matrix P . In all states, the probability of remaining in the same state is p and the probability of switching to another state is $1 - p$. The probability of switching to another state is divided equally between all other states, so P is a symmetric matrix with p in each element on the leading diagonal and $(1 - p)/(n - 1)$ in the off-diagonals.

3.2 REE with inflation sometimes outside the target range

We begin our analysis by looking at Rational Expectations Equilibria (REE) in which inflation is sometimes below, sometimes within, and sometimes above its target range. To that end, we assume there are 6 states in the exogenous Markov chain for ϵ_t , and impose restrictions on π^* so equilibrium inflation is below the target range when $\epsilon_t \in (-2, -1)$, in the range when $\epsilon_t \in (-1/3, 1/3)$, and above when $\epsilon_t \in (1, 2)$.⁵ The disturbance term is the only state variable in the model, so, abstracting for now from sunspot solutions, these REE are completely characterised by the level of inflation in each state of ϵ_t . Appendix A.1 uses a guess-and-verify approach to obtain the unique such REE (A.8).

$$\pi_t = \begin{cases} \frac{5(\epsilon_t - \alpha\sigma\kappa\pi^*)}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6 + 5\alpha\sigma\kappa} & \text{if } \epsilon_t \in (-2, -1) \\ \frac{5\epsilon_t}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6} & \text{if } \epsilon_t \in (-1/3, 1/3) \\ \frac{5(\epsilon_t + \alpha\sigma\kappa\pi^*)}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6 + 5\alpha\sigma\kappa} & \text{if } \epsilon_t \in (1, 2) \end{cases} \quad (9)$$

The nature of this REE depends on the degree of persistence in the Markov chain for the disturbance term ϵ_t . If there is only mild persistence then the term is inflationary when inflation is both inside and outside its target range, and the REE in Equation (9) is monotonic in ϵ_t . However, if the persistence is sufficiently high then the disturbance term is inflationary outside the target range but disinflationary inside it, which makes

⁴We work with a static IS curve to have a system that has a graphical representation. More generally, neural networks can predict vectors of outcomes and there is no need to restrict the number of variables that agents are forecasting. For example, in Section 6 the household-producer makes two forecasts when choosing consumption and setting prices.

⁵The restrictions on π^* needed to support the REE are that the target range is sufficient to cover inflation in two states but small enough that inflation is below or above the range in the other states.

the REE non-monotonic. A critical probability, $p^* = (6 + \sigma\kappa)/(6(1 + \sigma\kappa))$, determines whether the disturbance term is inflationary or deflationary inside the target range for inflation. Figure 2 sketches the REE for a mild degree of persistence $p < p^*$ that maintains monotonicity throughout, and a high degree of persistence $p > p^*$ where the disturbance term is disinflationary within the target range.⁶




Figure 2: Rational Expectations Equilibrium

The description of REE is completed by forward-looking expectations of inflation:

$$E_t \pi_{t+1} = \begin{cases} -\frac{(6p-1)(\epsilon_t - \alpha\sigma\kappa\pi^*)}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6 + 5\alpha\sigma\kappa} & \text{if } \epsilon_t \in (-2, -1) \\ -\frac{(6p-1)\epsilon_t}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6} & \text{if } \epsilon_t \in (-1/3, 1/3) \\ -\frac{(6p-1)(\epsilon_t + \alpha\sigma\kappa\pi^*)}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6 + 5\alpha\sigma\kappa} & \text{if } \epsilon_t \in (1, 2) \end{cases} \quad (10)$$

3.3 Learning

We assume that agents use neural network learning to form expectations, with a specification that is flexible enough for them to capture any nonlinearities in the economy. In our example with only two breakpoints, it is sufficient to have a neural network with one node that is always activated and two nodes that are activated by ReLU functions. The neural network then maps inputs $(1, \epsilon_t)$ to predicted inflation by:

$$E_t \pi_{t+1} = b_{21}(a_1 + b_{11}\epsilon_t) + b_{22}\max(a_2 + b_{12}\epsilon_t, 0) + b_{23}\max(a_3 + b_{13}\epsilon_t, 0), \quad (11)$$

where $w_1 \equiv (a_1, a_2, a_3, b_{11}, b_{12}, b_{13})$ and $w_2 \equiv (b_{21}, b_{22}, b_{23})$ are the parameters in the input and output layers. As it stands, the parameters in Equation (11) are not identified because we can multiply the input layer parameters by a positive constant and divide the output layer parameters by the same constant without changing the network's prediction. Without loss of generality, we rectify the issue by normalising the input layer coefficients

⁶Highly persistent shocks are disinflationary because the nominal interest rate does not react when inflation is within the target range. If ϵ_t is highly persistent then $E_t \epsilon_{t+1} \approx \epsilon_t$ and $E_t \pi_{t+1} \approx \pi_t$, which, in the absence of a reaction of R_t in the target range, implies $\pi_t \approx -\epsilon_t/(\sigma\kappa)$ by Equations (6)-(8), and a positive ϵ_t is associated with a negative π_t .

to $b_{11} \equiv b_{12} \equiv b_{13} \equiv 1$. This leaves six parameters for agents to learn.⁷

The neural network supports the REE for the parameter values $w_1^* \equiv (a_1^*, a_2^*, a_3^*)$ and $w_2^* \equiv (b_{21}^*, b_{22}^*, b_{23}^*)$ derived in Appendix A.2. With these values, the first node in Equation (11) is always activated, the second node is activated for $\epsilon_t \in (-1/3, 1/3, 1, 2)$, and the third node is activated for $\epsilon_t \in (1, 2)$. This is consistent with the REE outcomes in Figure 2, expectations being formed according to Equation (10), and is an application of the Universal Approximation Theorem discussed in Section 2.1.⁸

Agents apply the method of gradient descent to learn the parameters of the neural network. In period t they make prediction error $\zeta_t = \pi_t - E_{t-1}\pi_t$ and update their parameter estimates according to the recursive algorithm in Equation (3). The weighting functions for updating w_1 and w_2 are:

$$X_t^T w_2^T \psi'(X_t w_1) = \begin{pmatrix} 1 \\ \epsilon_{t-1} \end{pmatrix} \cdot \begin{pmatrix} b_{21} & b_{22} & b_{23} \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & \mathbb{I}_2 & 0 \\ 0 & 0 & \mathbb{I}_3 \end{pmatrix}, \quad (12)$$

$$\psi(X_t w_1)^T = \begin{pmatrix} a_1 + \epsilon_t & \max(a_2 + \epsilon_t, 0) & \max(a_3 + \epsilon_t, 0) \end{pmatrix}, \quad (13)$$

where $\mathbb{I}_2$ and $\mathbb{I}_3$ are indicator functions for when the second and third nodes are activated. If the activation functions are ReLU then the indicator functions are one when the node is activated and zero otherwise, unless the input to the function is exactly at a breakpoint where the derivative $\psi'(X_t w_1)$ is undefined. The possibility that derivatives may be undefined is not an issue in our example economy as the REE values of inflation are not at breakpoints in Figure 2, which ensures that derivatives are well-defined when parameter estimates are in the neighbourhood of their REE values. The need for derivatives to be well-defined in other environments can easily be satisfied by replacing the ReLU functions with suitably-convex sigmoid functions.

3.4 Learnability

We ask whether agents can learn the parameters of the neural network for the REE with inflation sometimes outside the target range, starting from initial estimates that are in the neighbourhood of their REE values. The answer is mostly no, because the REE is only learnable if the disturbance term is mildly persistent and interest rates react strongly when inflation is outside its target range. The result flows from the stability properties of the ordinary differential equation associated with neural network learning, following the procedure outlined in Section 2.3.

The starting point is the set of prediction errors that agents make when forming expectations using a neural network with fixed parameter estimates. Prediction errors depend on the difference between π_t and $E_{t-1}\pi_t$, which are functions of ϵ_t and ϵ_{t-1} in the example economy. With six states in the exogenous Markov chain for ϵ_t , the set of possible prediction errors has $6^2 = 36$ members. The procedure weights the set of prediction errors

⁷Software packages for fitting neural networks automatically normalise parameter estimates so there is no risk of underidentification in the numerical examples in Sections 4 and 6, even though they allow for a large number of nodes.

⁸The Universal Approximation Theorem (Cybenko, 1989) implies that any REE has a neural network representation, although it is silent on whether the neural network representation is learnable.

by Equations (12) and (13) and the unconditional probabilities of the states from the stationary distribution of the Markov chain. The result is a 6×1 vector function $h(\theta)$ that depends only on parameter estimates w_1, w_2 and structural parameters $\alpha, \sigma\kappa, \pi^*$.⁹ It shows the expected direction of updates when estimates are in the neighborhood of their REE values, with $h(\theta) = 0$ at the REE parameter values.

Agents can learn the parameters in their neural network if updating leads their estimates back to the REE values. Figure 3 uses $h(\theta)$ and the ordinary differential equation (4) to show the path that parameter estimates are expected to take after a perturbation of a_2 from its REE value, for $p = 1/4, \alpha = 5, \pi^* = 1, \sigma\kappa = 1$ and parameter estimates other than a_2 initialised at their REE values. As a_2 increases, agents make prediction errors that lead to them updating all their estimates, before the updates are unwound as estimates return to their REE values. Similar stability is observed when other parameter estimates are perturbed.




Figure 3: Responses to a perturbation to a_2 with ϵ_t mildly persistent.

The general condition for agents to learn the parameters of their neural network is that all the eigenvalues of the Jacobian of $h(\theta)$ have negative real parts when evaluated at the REE parameter values. $h(\theta)$ is a 6×1 vector so its Jacobian $Dh(\theta)$ is a 6×6 matrix, with each element a nonlinear function of the REE parameter values and structural parameters $p, \alpha, \pi^*, \sigma\kappa$. Rather than working with the eigenvalues directly, we note that the determinant of the Jacobian must be positive if all six eigenvalues are to be negative. Appendix A.5 derives two necessary conditions (14) and (15) for the determinant to be positive in the example economy, and shows that these are also sufficient for the Jacobian to have six negative eigenvalues.

$$p < \frac{6 + \sigma\kappa}{6(1 + \sigma\kappa)} = p^*, \quad (14)$$

$$p < \frac{4\alpha - 6 - \sigma\kappa(1 + \alpha)}{(4\alpha - 6)(1 + \sigma\kappa)}. \quad (15)$$

The first condition requires the probability of the disturbance term remaining in the same state to be less than p^* , the critical value that determines whether the disturbance term is inflationary or disinflationary when inflation is within its target range. It means the configuration with mild persistence in Figure 2a could be learnable but that with high persistence in Figure 2b is not. The second constraint further restricts the

⁹The detailed calculations are in Appendices A.3 and A.4, where we continue to abstract from sunspot solutions.

probability, so not every REE with mild persistence is learnable.

To understand the conditions, consider again a perturbation to a_2 that increases inflation expectations when inflation is within or above its target range. This raises inflation more than proportionately as long as inflation stays in the target range, leading to negative prediction errors and updates to a_2 that take it even further away from its REE value. The estimate of a_2 diverges until either inflation leaves the target range or the disturbance term switches to a state where the nominal interest rate reacts to inflation, at which point prediction errors turn positive and updates to a_2 gradually bring it back to its REE value. If the probability of remaining in the same state is small enough then the convergent dynamics outside the target range dominate and the REE is learnable, as in Figure 3. Otherwise, it is not and inflation is forced out of its target range.

3.5 Sunspot fluctuations


The example nonlinear economy permits sunspot fluctuations in REE under certain conditions. Appendix B.1 has an illustrative case in which an exogenous non-fundamental sunspot process independently switches between two states. The key condition is a ‘resonance frequency’ requirement that the probability of the sunspot process switching state is compatible with self-fulfilling fluctuations. For a sunspot to survive in REE, it needs to induce changes in expectations and outcomes that are mutually consistent. The sunspot is attached to $E_t\pi_{t+1}$ when $\epsilon_t \in (-1/3, 1/3)$, and passes through to π_t with a coefficient of $1 + \sigma\kappa$ via the New Keynesian Phillips curve (6) and IS curve (7). These movements can only be compatible if the sunspot switches state with a specific probability. The requirement that the probability lies between zero and one imposes a further condition, that $\sigma\kappa$ is large enough relative to p .

The REE with sunspot fluctuations is compatible with an expanded neural network that includes the sunspot term in the node that is always activated. The details are in Appendix B.2, where a new parameter on the sunspot term increases the number of parameters that agents have to learn to seven. Learnability then depends on all the eigenvalues of the 7×7 matrix of the Jacobian of $h(\theta)$ having negative real parts. This is never the case in the example nonlinear economy, because the eigenvalue associated with learning the sunspot parameter is zero whenever the ‘resonance frequency’ condition needed to support sunspot fluctuations is satisfied. Intuitively, for a sunspot to survive it must be that non-fundamental changes in expectations are self-fulfilling, but the conditions needed for that to occur mean that changes in the sunspot parameter are equally self-fulfilling. Ironically, the ‘resonance frequency’ condition that facilitates sunspot fluctuations in REE is precisely the condition that precludes agents and the neural network from learning them.

3.6 REE with inflation always outside the target range

High persistence makes it impossible for agents to learn the REE with inflation sometimes below, within, and above its target range. This raises the question of what happens to inflation when persistence is high. We know that learning dynamics are unstable and inflation is driven outside the target range, but what happens then? The answer is that parameter estimates converge to values consistent with a learnable REE in which inflation is always outside the target range. A learnable REE of this type always exists when persistence is

Figure 4 illustrates the detailed calculations in Appendix C. The learnable REE outside the target range is in red, with inflation below or above the range depending on whether the disturbance term is negative or positive. In black is the non-learnable REE where inflation is sometimes below, within, and above the target range. The REE are the same for $\epsilon_t \in (-2, -1, 1, 2)$ but differ for $\epsilon_t \in (-1/3, 1/3)$. Dynamics follow the arrows, as learning drives parameter estimates and inflation outside the target range to the learnable REE.



4 Numerical example with target range for inflation

4.1 Environment

11

ϕ_π and α imply that policy only satisfies the Taylor Principle when inflation is outside its target range.

$$\pi_t = \beta E_t \pi_{t+1} + \kappa y_t + \epsilon_{\pi,t} \quad (16)$$

$$y_t = \eta y_{t-1} - \sigma(R_t - E_t \pi_{t+1}) + \epsilon_{y,t} \quad (17)$$

$$R_t = \begin{cases} \phi_\pi \pi_t + \alpha(\pi_t + \pi^*) & \text{if } \pi_t < -\pi^* \\ \phi_\pi \pi_t & \text{if } -\pi^* \leq \pi_t \leq \pi^* \\ \phi_\pi \pi_t + \alpha(\pi_t - \pi^*) & \text{if } \pi_t > \pi^* \end{cases} \quad (18)$$

The disturbance terms $(\epsilon_{\pi,t}, \epsilon_{y,t})$ are independent AR(1) processes with persistence parameters (ρ_π, ρ_y) and normally-distributed innovations of standard deviations (σ_π, σ_y) . The baseline parameter values for our numerical exercises are in Table 1. The numerical example reduces to the analytic example in the previous section if $\epsilon_{\pi,t}$ is discrete and $\phi_\pi = \eta = \sigma_y = 0$.

β	κ	η	σ	ϕ_π	α	ρ_π	ρ_y	σ_π	σ_y
0.95	0.05	(0.75, 0.95)	0.25	0.5	0.75	0.5	0.5	0.2	0.2

Table 1: Baseline parameter values

Agents use neural network learning to form predictions of inflation $E_t \pi_{t+1}$ in the next period, with the vector of inputs $(\epsilon_{\pi,t}, \epsilon_{y,t}, y_{t-1})$. There is no cost to redundancy in the network other than computation time, so we set the number of nodes equal to $K = 32$. We generate numerical results by recursively simulating the model and retraining the neural network until its parameters converge. Starting from some initialisation of the neural network, we repeatedly produce long samples of the exogenous and endogenous variables and use the simulated data to update the neural network's parameters.

4.2 Deterministic steady states

The deterministic steady states of the model occur at the intersection of steady-state versions of the New Keynesian Phillip curve, IS curve and policy rule. Figure 5 shows the steady-state Phillips curve as the straight dashed line; the steady-state IS curve and policy rule are combined in the solid z-shape (details in Appendix D.1). The number of deterministic steady states depends on whether IS cuts PC from below or above when inflation is zero, which in turn depends on whether:

$$\eta \begin{matrix} \leq \\ > \end{matrix} 1 - \frac{\sigma \kappa (1 - \phi_\pi)}{1 - \beta}. \quad (19)$$

There is one deterministic steady state when the coefficient η on lagged output in the Phillips curve is small, but three when it is large. Multiplicity occurs because inertia in the IS curve flattens the steady-state IS curve and policy rule relative to the steady-state Phillips curve. In Appendix D.2, we show that the model is locally determinate in the neighbourhood of the (0,0) deterministic steady state when inertia is low, and locally indeterminate when inertia is high. When high inertia produces additional deterministic steady states as in Figure 5b, they are always locally determinate.




Figure 5: Deterministic steady states

4.3 REE

The fixed point to which neural network learning converges is not necessarily an REE unless it leads to agents forming expectations that are rational. We test for this rationality using the accuracy test of Den Haan and Marcet (1994), which checks that agents' expectational errors are unpredictable based on the information available at the time expectations are formed. The test asks whether $\mathbb{E}[e_{t+1|\theta} \otimes h(s_t)] = 0$, i.e., whether the vector of expectations errors $e_{t+1|\theta}$ generated under beliefs θ is orthogonal to a function $h(s_t)$ of an information set s_t . The test statistic is calculated from the properties of expectations errors in simulations of the model, and has a theoretical distribution under the null hypothesis of unpredictability that has 5% of draws in its upper and lower tails. We make 1,000 draws of 500 periods to obtain 1,000 test statistics, accepting the null hypothesis that expectations are rational if close to 5% of the test statistics fall in each of the upper and lower tails. The test is a demanding standard yet the neural network passes it comfortably, giving us confidence that our solution has the properties of an REE.

4.4 Equilibrium dynamics

The phase diagrams of the system are in Figure 6. The two steady-state relationships are as before, so the system still has one deterministic steady state when inertia is low and three when it is high. The dynamics with low inertia are convergent but the dynamics with high inertia diverge from the (0,0) deterministic steady state. This is consistent with the (0,0) deterministic steady state being locally determinate when η is small and locally indeterminate when η is large, and the Ellison and Pearlman (2011) assertion that locally determinate REE are always learnable but locally indeterminate REE are difficult to learn. When inertia is high, the deterministic steady state at (0,0) becomes a source that repels the system along a unique path towards one of the stable outer steady states.




Figure 6: Equilibrium dynamics under neural network learning

The convergent dynamics in Figure 6a imply that inflation reverts to the middle of its target range when η is low. In simulations, inflation and output are centred on zero and, at least in the baseline parameterisation, inflation rarely strays outside its target range. We found no evidence that neural network learning supports sunspot fluctuations. These results mirror those in the analytic example with low persistence in the disturbance term, where the only learnable REE has no sunspot fluctuations and inflation is often within its target range. The divergent dynamics in Figure 6b mean it is rare for inflation to be in the target range when η is high. Inflation spends most of its time in the neighbourhood of one of the outer steady states, and periods within its target range are typically short-lived as output and inflation are pushed towards one of the outer steady states. The parallel in the analytic example is when the disturbance term is highly persistent and inflation is never in its target range in the learnable REE.

The equilibrium paths in Figure 6b lead to rest points that are not exactly centred on the deterministic steady states. This is consistent with the literature on “deflationary bias” in nonlinear stochastic RE models with a zero lower bound constraint (Nakata and Schmidt, 2019). When inflation is above its target range, the possibility that future shocks may push inflation into the target range gives forward-looking agents an incentive to reduce their prices, biasing inflation down from its deterministic steady-state value. Output is above the deterministic steady state because monetary policy accommodates the deflationary bias (Bianchi et al., 2021). An equal and opposite “inflationary bias” explains the location of the rest point relative to the deterministic steady state when inflation is below its target range.

4.5 Perfect foresight

The dynamics uncovered by neural network learning are not the only REE of the system. One class of REE that is readily calculable is based on perturbations of perfect foresight paths. To describe the dynamics of inflation under perfect foresight, we recognise that $E_t \pi_{t+1} = \pi_{t+1}$ and abstract from the disturbance terms to express π_t and y_t in terms of π_{t-1} and y_{t-1} . These dynamics are shown by the phase diagrams in Figure 7. The red arrows represent perfect foresight paths along which the system will evolve from a given starting

point in the (y_{t-1}, π_{t-1}) space of state and control variables.

In Figure 7a, there is a unique stable perfect foresight path that resembles the path under neural network learning in Figure 6a. It follows that simulations based on perturbations of perfect foresight paths are likely to be a good approximation of the learnable REE when inertia is low. In contrast, the perfect foresight paths in Figure 7b are very different from the equilibrium dynamics with neural network learning in Figure 6b. When inertia is high, there are multiple perfect foresight paths in the neighbourhood of the $(0,0)$ deterministic steady state and unique stable perfect foresight paths that converge to each of the outer steady states. A researcher approximating the behaviour of inflation as perturbations around perfect foresight paths is in danger of grievously misidentifying equilibrium dynamics. The analogy in the analytic example would be mistakenly using the non-learnable REE to characterise dynamics when the disturbance term is highly persistent.




Figure 7: Perfect foresight paths

5 Empirical tests for indeterminacy

This Section applies the Lubik and Schorfheide (2004) test to simulated data from our model. The test compares the probability that the data is generated by a locally-determinate system as opposed to being perturbations of perfect foresight paths in a locally-indeterminate system. Since learnable REE dynamics resemble neither a saddlepath stable system nor perturbations to perfect foresight paths when inertia is high, we investigate whether the test may misleadingly favour sunspots as a driver of the data.

5.1 Introducing sunspots

The Lubik and Schorfheide (2004) test is based on estimating the log-linearised dynamics around a steady state of a model. Without loss of generality, the log-linearised model can be expressed as:

$$\mathbb{E}_t \hat{x}_{t+1} = A \hat{x}_t + B \epsilon_t, \quad (20)$$

and the determinacy of the system depends on the eigenvalues of the matrix A .¹⁰ In our model with one control and one state variable, the system is determinate if only one of the eigenvalues is inside the unit circle. If both are inside then the system is indeterminate. The model is solved under rational expectations by adding and subtracting $A\mathbb{E}_{t-1}\hat{x}_t$ from the right side of (20) and applying the Jordan decomposition $A = P^{-1}\Lambda P$:

$$PE_t\hat{x}_{t+1} = \Lambda P\mathbb{E}_{t-1}\hat{x}_t + \Lambda P(\hat{x}_t - \mathbb{E}_{t-1}\hat{x}_t) + PB\epsilon_t. \quad (21)$$

The diagonality of Λ facilitates a decoupling of the system with $\tilde{x}_t = P\hat{x}_t$ and $\tilde{\epsilon}_t = PB\epsilon_t$:

$$\mathbb{E}_t\tilde{x}_{1,t+1} = \Lambda_1\mathbb{E}_{t-1}\tilde{x}_{1,t} + \Lambda_1(\tilde{x}_{1,t} - \mathbb{E}_{t-1}\tilde{x}_{1,t}) + \tilde{\epsilon}_{1,t}, \quad (22)$$

$$\mathbb{E}_t\tilde{x}_{2,t+1} = \Lambda_2\mathbb{E}_{t-1}\tilde{x}_{2,t} + \Lambda_2(\tilde{x}_{2,t} - \mathbb{E}_{t-1}\tilde{x}_{2,t}) + \tilde{\epsilon}_{2,t}. \quad (23)$$

If the system is determinate then one of Λ_1 or Λ_2 will be outside the unit circle and the decoupled equation associated with it has a single stable solution $\tilde{x}_{i,t} = 0 \forall t$. This condition identifies the saddlepath relationship between the state and control variables, and guarantees that (22) and (23) have a unique solution that can be estimated by standard techniques.

If Λ_1 and Λ_2 are both inside the unit circle then the system is indeterminate and (22) and (23) do not have a unique solution. In this case, Farmer et al. (2015) suggest closing the model with a sunspot process that selects between the multiple perfect foresight paths that converge to the steady state. The sunspot process ζ_t has to respect rational expectations to the extent that $\zeta_t = \pi_t - \mathbb{E}_{t-1}\pi_t$, but can be correlated with innovations to the disturbance terms:¹¹

$$\mathbb{E}_{t-1} \begin{pmatrix} \epsilon_{\pi,t} \\ \epsilon_{y,t} \\ \zeta_t \end{pmatrix} \begin{pmatrix} \epsilon_{\pi,t} \\ \epsilon_{y,t} \\ \zeta_t \end{pmatrix}' = \begin{pmatrix} \sigma_\pi & 0 & \omega_{\pi,\zeta} \\ 0 & \sigma_y & \omega_{y,\zeta} \\ \omega_{\pi,\zeta} & \omega_{y,\zeta} & \sigma_\zeta \end{pmatrix} \quad (24)$$

The requirement that sunspots satisfy rational expectations implies there is a unique solution to (22) and (23), where the sunspot is consistent with the variance-covariance matrix (24). Standard techniques can then be employed to estimate $(\omega_{\pi,\zeta}, \omega_{y,\zeta}, \sigma_\zeta)$ and other parameters.

5.2 Estimation results

The Lubik and Schorfheide (2004) test for indeterminacy amounts to estimating separate linear models with and without sunspots on the same data and comparing their fit via a posterior odds ratio. The models are estimated in Dynare using Metropolis-Hastings MCMC, as described in Guerrón-Quintana and Nason (2013). We perform the test on 10,000 periods of simulated data from each of our learnable REEs.

¹⁰Blanchard and Kahn (1980) or Sims (2002).

¹¹Lubik and Schorfheide (2004) explicitly model a dependence of the sunspot on fundamental shocks.

Parameter	Prior mean	Posterior mean for $\eta = 0.75$		Posterior mean for $\eta = 0.95$	
		Determinate	Indeterminate	Determinate	Indeterminate
β	$\mathcal{N}(1, 1)$	0.06 (-0.57, 0.84)	1.90 (1.77, 2.00)	0.40 (0.26, 0.47)	1.37 (1.20, 1.57)
κ	$\mathcal{N}(0, 1)$	0.16 (0.07, 0.25)	-0.10 (-0.14, -0.05)	0.48 (0.35, 0.58)	-0.26 (-0.42, 0.14)
η	$\mathcal{N}(1, 1)$	0.75 (0.73, 0.78)	0.75 (0.73, 0.77)	0.98 (0.97, 0.99)	0.60 (0.49, 0.70)
σ	$\mathcal{N}(0, 1)$	-0.09 (-0.22, 0.00)	0.15 (0.01, 0.36)	0.001 (-0.01, 0.02)	0.65 (0.52, 0.80)
ϕ_π	$\mathcal{N}(2, 1)$	1.32 (0.62, 1.94)	0.68 (0.38, 0.99)	1.81 (1.64, 1.98)	0.22 (0.08, 0.39)
ρ_π	$\mathcal{N}(0, 1)$	0.51 (0.49, 0.52)	0.24 (-0.05, 0.54)	0.50 (0.48, 0.52)	0.42 (0.38, 0.47)
ρ_y	$\mathcal{N}(0, 1)$	0.49 (0.45, 0.52)	0.50 (0.47, 0.53)	0.48 (0.47, 0.49)	0.51 (0.48, 0.55)
σ_π	$\mathcal{IG}(0.5, 2)$	0.37 (0.22, 0.49)	0.31 (0.17, 0.46)	0.31 (0.28, 0.34)	0.24 (0.21, 0.27)
σ_y	$\mathcal{IG}(0.5, 2)$	0.21 (0.20, 0.21)	0.21 (0.20, 0.21)	0.21 (0.21, 0.21)	0.14 (0.12, 0.15)
σ_ζ	$\mathcal{IG}(0.5, 2)$	-	0.34 (0.31, 0.38)	-	0.41 (0.40, 0.43)
$\omega_{\pi, \zeta}$	$\mathcal{B}(0, 0.3, -1, 1)$	-	· (-, +)	-	0.39 (0.19, 0.58)
$\omega_{y, \zeta}$	$\mathcal{B}(0, 0.3, -1, 1)$	-	· (-, +)	-	-0.10 (-0.46, 0.28)
Log data density		-3012	-3019	-3499	-3496

Table 2: Bayesian estimation results

The Bayesian estimation results are in Table 2, where the posteriors from draws of the parameters that imply determinacy and indeterminacy are reported in the third to sixth columns. The results show support for the determinate model when η is low but a preference for indeterminacy when η is high. In the latter case, the log posterior for the determinate model is -3499 whereas for the indeterminate model it is -3496, so the posterior odds ratio favours indeterminacy and the simulated data prefers the sunspot model over the determinate linear rational expectations model. The estimated sunspot process is positively correlated with both innovations to the disturbance terms, although it is noticeably more volatile than either of them.

6 Numerical example with a Zero Lower Bound

Another much-studied nonlinear economy is one in which the nominal interest rate is subject to a zero lower bound. Since Benhabib et al. (2001), it is known that these economies may possess a second, deflationary, steady state. Policymakers take such a possibility seriously, worrying that negative shocks “could push the economy into an unintended, low nominal interest rate steady state” (Bullard, 2010). This Section draws on our experience with neural network learning to discuss the likely implications for economic dynamics.

6.1 Discussion

The zero lower bound creates a nonlinearity similar to those in our examples with a target range for inflation. When inflation is above a certain threshold, monetary policy is free to raise the nominal interest rate and bring inflation under control. When below, the nominal interest rate is constrained at zero and monetary policy cannot be used to stimulate the economy. In Benhabib et al. (2001), this leads to two steady states that are differentiated by whether or not the zero lower bound is binding. Inflation is typically lower and may even be deflationary in the steady state where policy is constrained. However, Christiano et al. (2018) claim this is not a concern because the constrained steady state is not learnable.

Our investigations into neural network learning confirm that imposing a zero lower bound on an otherwise standard New Keynesian model renders REE dynamics unlearnable.¹² Equilibrium dynamics are well-defined when conditioned on the estimated parameters of agents’ neural networks, but successive updating of those parameters by the method of gradient descent leads to increasingly pessimistic beliefs whenever there is a possibility that the zero lower bound will bind in the future. Richter and Throckmorton (2015) reach a similar conclusion, with numerical solution methods not converging if zero lower bound episodes are frequent and long-lasting. The dynamics under successive rounds of neural network learning resemble the unbounded deflationary spirals under adaptive learning of Evans et al. (2008).

Macroeconomists have struggled with the non-learnability of REE in models with a zero lower bound. One approach is to relinquish full rationality and solve forward under a perfect foresight assumption that future shocks are zero and the economy will be at the “target rate” in some future period (Braun and Körber, 2011; Christiano et al., 2018; Guerrieri and Iacoviello, 2015).¹³ Another is to augment the model with a mechanism that generates a third steady state at which there is secular stagnation. Eggertsson et al. (2019) do this via long-run non-neutralities in aggregate supply and an OLG structure that inverts the slope of the aggregate demand curve when the economy is in a liquidity trap at the zero lower bound. Evans et al. (2022) instead combine supply non-neutralities with government spending that puts a floor under aggregate demand.

None of the papers discussed above identify a fully stochastic nonlinear learnable REE with a zero lower bound on the nominal interest rate, a deficiency we turn to by introducing neural network learning into the model of Evans et al. (2022). Our version of their model adds consumption habits and a rule that endogenises government spending. In what follows we sketch the key equations, relegating details to Appendix E.

¹²We also find no evidence of learnable REE sunspot solutions of the type in Aruoba et al. (2018).

¹³The risks involved with such perfect foresight methods were highlighted in Section 4.5.

6.2 Environment

There is a continuum of identical household-producers indexed by i that maximise utility:

$$\mathbb{E}_{0,i} \sum_{t=0}^{\infty} \beta^t \left\{ \log(c_{t,i} + \xi_t g_t - \lambda c_{t-1}) + \chi \log \left(\frac{M_{t-1,i}}{P_t} \right) - (1 + \epsilon)^{-1} h_{t,i}^{1+\epsilon} - \Phi \left(\frac{P_{t,i}}{P_{t-1,i}} \right) \right\}, \quad (25)$$

where $c_{t,i}$ is the consumption aggregator consumed by i , g_t is exogenous government spending per capita, and the term λc_{t-1} in lagged aggregate consumption is the external habit. $M_{t,i}$ denotes nominal money balances, P_t is the aggregate price level, and $h_{t,i}$ is the labour input into production of consumption good i . $P_{t,i}$ is the price of consumption good i , and $\Phi(\cdot) \geq 0$ is a convex pricing friction with $\Phi = 0$ at the inflation target π^* .

The household-producers are subject to a production function, $y_{t,i} = A_t h_{t,i}^\alpha$, a non-negativity constraint on consumption, $c_{t,i} \geq 0$, and an intertemporal budget constraint:

$$c_{t,i} + m_{t,i} + b_{t,i} + \Upsilon_{t,i} = \frac{m_{t-1,i}}{\pi_t} + R_{t-1} \frac{b_{t-1,i}}{\pi_t} + \frac{P_{t,i}}{P_t} y_{t,i}, \quad (26)$$

where $m_t = M_t/P_t$ denotes real money balances, b_t is the real quantity of risk-free one-period nominal bond holdings at the end of period t , and $\Upsilon_{t,i}$ is lump-sum taxes collected by the government. $\pi_t = P_t/P_{t-1}$ is the inflation rate, R_{t-1} is the nominal interest rate between $t-1$ and t , and $y_{t,i}$ is output of good i . Each household-firm faces a downward-sloping demand curve:

$$P_{t,i} = \left(\frac{y_{t,i}}{y_t} \right)^{-1/\nu} P_t, \text{ where } P_t = \left[\int_0^1 P_{t,i}^{1-\nu} \right]^{1/(1-\nu)}, \quad (27)$$

with elasticity of substitution between consumption goods $\nu > 1$. Our model has two exogenous processes: a productivity shock A_t and a preference shock ξ_t . The market clearing conditions are $y_{t,i} = y_t = c_t + g_t$.

6.3 Policy

The nominal interest rate R_t reacts to inflation and output, subject to the zero lower bound constraint:

$$R_t = \max \left\{ R^* \left(\frac{\pi_t}{\pi^*} \right)^{\phi_\pi} \left(\frac{y_t}{y^*} \right)^{\phi_y}, 1 \right\}, \quad (28)$$

with targets (R^*, π^*, y^*) that are consistent in steady state. Government spending g_t reacts to inflation:

$$g_t = \frac{\bar{g}}{1 + e^{k(\pi_t - \pi^*)}}. \quad (29)$$

Under this rule, spending is $\bar{g}/2$ when inflation is at target and increases as inflation falls below target, with the parameter k determining the speed at which spending increases to its limit $\bar{g}$.¹⁴ The government sets lump-sum taxes to balance their period budget constraint: $b_t + m_t + \Upsilon_t = g_t + (m_{t-1} + R_{t-1}b_{t-1})/\pi_t$, which the household-producers internalise.

¹⁴If $k = 0$ then government spending is constant, as in Evans et al. (2022). We endogenise government spending to avoid an absolute lower limit on output, which makes it difficult for household-producers to learn any REE with a neural network.

6.4 Equilibrium conditions

Optimal consumption of the household-producer maximises utility (25), subject to the flow budget constraint (26) and the non-negativity constraint on consumption. They are on their consumption Euler equation (31) when desired consumption c_t^* is positive, otherwise their consumption is zero. Household-producers still benefit from government spending when c_t is zero, so their marginal utility of consumption remains bounded.

$$c_t = \max(c_t^*, 0), \quad (30)$$

where c_t^* satisfies

$$\frac{1}{c_t^* + \xi_t g_t - \lambda c_{t-1}} = \beta R_t E_t \left(\frac{\pi_{t+1}^{-1}}{c_{t+1}^* + \xi_{t+1} g_{t+1} - \lambda c_t} \right). \quad (31)$$

Optimal price setting by the household-producer maximises utility (25), subject to the flow budget constraint (26) and the downward-sloping demand curve (27). It implies a New Keynesian Phillips Curve (32), where the “output gap” terms are the marginal disutility of labour input and the marginal utility of consumption associated with changing prices. The convex friction $\Phi(\cdot)$ makes price setting forward-looking.

$$\Phi'(\pi_t)\pi_t = \frac{\nu}{\alpha} \left(\frac{c_t + g_t}{A_t} \right)^{\frac{\epsilon+1}{\alpha}} + \frac{1-\nu}{c_t + \xi_t g_t - \lambda c_{t-1}} (c_t + g_t) + \beta E_t \Phi'(\pi_{t+1})\pi_{t+1} \quad (32)$$

The Euler equation for consumption (30)-(31) and the New Keynesian Phillips Curve (32) combine with market clearing $y_t = c_t + g_t$ and the policy rules for the nominal interest rate (28) and government spending (29) to determine the model’s endogenous variables (c_t, g_t, π_t, R_t) . The technology and preference shocks follow independent exogenous AR(1) processes with Gaussian innovations. Details are in Appendix E.1.

6.5 Deterministic steady states

The model’s equilibrium reduces to two relationships in output and inflation, derived from the consumption Euler equation and the New Keynesian Phillips Curve by substituting in for the nominal interest rate and government spending using the policy rules. Appendix E.2 solves for deterministic steady states. The steady-state Euler equation is downward-sloping when $\pi > \beta$ because then consumption is positive, the zero lower bound does not bind, and multiple combinations of (π, y) in the policy rule for the nominal interest rate are consistent with the consumption Euler equation. It is horizontal when $\pi = \beta$ as any level of steady-state output in the interval $[\underline{y}, \bar{y}]$ is compatible with steady-state inflation being equal to β . When $\pi < \beta$, the steady-state Euler equation is again downward-sloping because consumption is constrained at zero, the zero lower bound binds, and output is determined solely by government spending that increases as inflation falls. The steady-state New Keynesian Phillips Curve is upward sloping due to the nonlinear pricing friction $\Phi(\cdot)$.

The steady-state Euler equation and Phillips curve are illustrated in Figure 8, for a Linex specification of the pricing friction and the parameter values in Appendix E.3. There are three deterministic steady states at which the steady-state relationships intersect. On the right is the “target” steady state where inflation, output, and the nominal interest rate are all at their respective targets. In the middle is the “unlearnable” steady state at which the zero lower bound binds, as in Benhabib et al. (2001). On the left is the steady state

at which consumption is constrained at zero and output is determined entirely by government spending, akin to the third “secular stagnation” equilibrium in Eggertsson et al. (2019).




Figure 8: Deterministic steady states and equilibrium dynamics in the numerical example with a ZLB

6.6 Equilibrium dynamics

In a learnable REE, the household-producer uses neural network learning to form the expectations they need to choose consumption and set prices optimally. For consumption, they need to forecast $\mu_{t+1}\pi_{t+1}^{-1}$ in the Euler equation (31), where μ_{t+1} is the marginal utility of consumption. For price setting, they need to forecast $\Phi'(\pi_{t+1})\pi_{t+1}$ in the Phillips Curve (32). The state variables are $(y_{t-1}, g_{t-1}, A_t, \xi_t)$. As in the numerical example in Section 4, we identify the learnable REE through simulation and repeated updating of the parameters in the household-producer’s neural network using the method of gradient descent. Consistent with the analytical and numerical results in the examples with a target range for inflation, in the learnable REE the economy spends most of its time near to either the target or secular stagnation steady states.

The learnable REE is superimposed on Figure 8 in blue. In the now familiar pattern, equilibrium dynamics diverge from the “unlearnable” deterministic steady state at which the zero lower bound constraint binds. Depending on shocks, the economy then spends extended periods near the “target” or “secular stagnation” deterministic steady states. As in the previous example (Figure 6b), a deflationary bias reduces the rest point of inflation when inflation is high and an inflationary bias has the opposite effect when inflation is low. The two biases are not symmetric in this environment, so do not cancel near the middle deterministic steady state.

7 Conclusion

We see neural network learning as a practical and powerful tool for the analysis of nonlinear macroeconomies. Any REE has a neural network representation, which can be checked for learnability either analytically or numerically. Neural network learning therefore offers both a solution method and an equilibrium selection device. In the analytic examples of Section 3, it explains why inflation may always be outside a target range if disturbances are persistent, and that REE with sunspot fluctuations are typically not learnable. The insights are confirmed in the numerical example of Section 4, where too much endogenous persistence drives inflation away from the target range. Inflation is similarly pushed out in Section 6 with a zero lower bound.

The dynamics of REE under neural network learning are similar, whether there is a target range for inflation or a zero lower bound on the nominal interest rate. If persistence is high then inflation spends most of its time outside the target range, only occasionally passing through on its way to the other side. In simulations of the zero lower bound model, inflation is rarely close to the middle deterministic steady state. To the unsuspecting eye, the dynamics of the model are suggestive of a Markov switching specification (Hamilton, 1989) or an environment with endogenous sunspots (Aruoba et al., 2017). The empirical tests for indeterminacy reported in Section 5 suggest this is not a trivial possibility, nor do simple (unreported) simulation-based tests that support a two-state Markov-switching specification. Neural network learning in a nonlinear environment may therefore explain and provide microfoundations for some popular empirical approaches.

That our REE with sunspot fluctuations are not learnable should be of comfort to policymakers or legislators contemplating the introduction of a target range for inflation or facing a zero lower bound on the nominal interest rate. It is unlikely that the economy will suffer from unwanted sunspot volatility if agents form expectations using neural networks, and the presence of multiple rest points suggests that a suitably-designed policy could anchor inflation close to its target steady in the zero lower bound example. Indeed, we obtain learnable REE with dynamics that always fluctuate around the same rest point if we reduce the volatility of disturbances in our numerical examples. Such benefits need to be weighed against a concern that agents may be unable to learn an REE where inflation spends prolonged periods in its target range.

Appendix for

Neural Network Learning for Nonlinear Economies

A Analytic example without sunspot fluctuations

A.1 REE

The initial guess is that inflation is in the target range only when $\epsilon_t \in (-1/3, 1/3)$. It is below the target range when $\epsilon_t \in (-2, -1)$ and above it when $\epsilon_t \in (1, 2)$. The second column of Table A.1 assigns undetermined coefficients $A \dots F$ to the REE value of inflation for each ϵ_t . The rows of the third column use the transition matrix of the Markov chain for ϵ_t to derive the REE prediction of inflation in the next period at each ϵ_t . The fourth column is the reaction of the nominal interest rate R_t , taken from Equation (8) under the assumption that the interest rate does not react when $\epsilon_t \in (-1/3, 1/3)$.

ϵ_t	π_t	$E_t \pi_{t+1}$	R_t
-2	A	$pA + ((1-p)/5)(B + C + D + E + F)$	$\alpha\sigma\kappa(A + \pi^*)$
-1	B	$pB + ((1-p)/5)(A + C + D + E + F)$	$\alpha\sigma\kappa(B + \pi^*)$
-1/3	C	$pC + ((1-p)/5)(A + B + D + E + F)$	0
1/3	D	$pD + ((1-p)/5)(A + B + C + E + F)$	0
1	E	$pE + ((1-p)/5)(A + B + C + D + F)$	$\alpha\sigma\kappa(E - \pi^*)$
2	F	$pF + ((1-p)/5)(A + B + C + D + E)$	$\alpha\sigma\kappa(F - \pi^*)$

Table A.1: Expectations and nominal interest rate in REE

Inflation and the output gap at each ϵ_t are determined by the New Keynesian Phillips curve (6), the IS curve (7), and the REE predictions of inflation and reactions of nominal interest rates in Table A.1.

$$A = (1 + \sigma\kappa) [pA + ((1-p)/5)(B + C + D + E + F)] - \alpha\sigma\kappa(A + \pi^*) - 2 \quad (\text{A.1})$$

$$B = (1 + \sigma\kappa) [pB + ((1-p)/5)(A + C + D + E + F)] - \alpha\sigma\kappa(B + \pi^*) - 1 \quad (\text{A.2})$$

$$C = (1 + \sigma\kappa) [pC + ((1-p)/5)(A + B + D + E + F)] - 1/3 \quad (\text{A.3})$$

$$D = (1 + \sigma\kappa) [pD + ((1-p)/5)(A + B + C + E + F)] + 1/3 \quad (\text{A.4})$$

$$E = (1 + \sigma\kappa) [pE + ((1-p)/5)(A + B + C + D + F)] - \alpha\sigma\kappa(E - \pi^*) + 1 \quad (\text{A.5})$$

$$F = (1 + \sigma\kappa) [pF + ((1-p)/5)(A + B + C + D + E)] - \alpha\sigma\kappa(F - \pi^*) + 2 \quad (\text{A.6})$$

Equations (A.1)-(A.6) are a system of six linear equations in six unknowns, $A \dots F$.¹⁵

$$\begin{aligned}
 A &= -\frac{5(2 + \alpha\sigma\kappa\pi^*)}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6 + 5\alpha\sigma\kappa} & B &= -\frac{5(1 + \alpha\sigma\kappa\pi^*)}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6 + 5\alpha\sigma\kappa} & C &= -\frac{5}{3(\sigma\kappa - 6p(1 + \sigma\kappa) + 6)} \\
 D &= \frac{5}{3(\sigma\kappa - 6p(1 + \sigma\kappa) + 6)} & E &= \frac{5(1 + \alpha\sigma\kappa\pi^*)}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6 + 5\alpha\sigma\kappa} & F &= \frac{5(2 + \alpha\sigma\kappa\pi^*)}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6 + 5\alpha\sigma\kappa}
 \end{aligned} \quad (\text{A.7})$$

¹⁵REE_6.sigma.m

The piecewise linearity of the solution (A.7) permits a compact representation of the REE:

$$\pi_t = \begin{cases} \frac{5(\epsilon_t - \alpha\sigma\kappa\pi^*)}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6 + 5\alpha\sigma\kappa} & \text{if } \epsilon_t \in (-2, -1) \\ \frac{5\epsilon_t}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6} & \text{if } \epsilon_t \in (-1/3, 1/3) \\ \frac{5(\epsilon_t + \alpha\sigma\kappa\pi^*)}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6 + 5\alpha\sigma\kappa} & \text{if } \epsilon_t \in (1, 2) \end{cases} \quad (\text{A.8})$$

$$E_t\pi_{t+1} = \begin{cases} -\frac{(6p-1)(\epsilon_t - \alpha\sigma\kappa\pi^*)}{6\sigma\kappa - 6p(1 + \sigma\kappa) + 6 + 5\alpha\sigma\kappa} & \text{if } \epsilon_t \in (-2, -1) \\ -\frac{(6p-1)\epsilon_t}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6} & \text{if } \epsilon_t \in (-1/3, 1/3) \\ -\frac{(6p-1)(\epsilon_t + \alpha\sigma\kappa\pi^*)}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6 + 5\alpha\sigma\kappa} & \text{if } \epsilon_t \in (1, 2) \end{cases} \quad (\text{A.9})$$

It remains to verify the initial guess that inflation is in the target range only when $\epsilon_t \in (-1/3, 1/3)$. It suffices to show that inflation is in the target range when $\epsilon_t = \pm 1/3$ and above when $\epsilon_t = \pm 1$, which occurs if:

$$\frac{5}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6} > \pi^* > \frac{5}{3(\sigma\kappa - 6p(1 + \sigma\kappa) + 6)} \quad \text{if } p < (6 + \sigma\kappa)/(6(1 + \sigma\kappa)) \quad (\text{A.10})$$

$$\frac{5}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6} < -\pi^* < \frac{5}{3(\sigma\kappa - 6p(1 + \sigma\kappa) + 6)} \quad \text{if } p > (6 + \sigma\kappa)/(6(1 + \sigma\kappa)) \quad (\text{A.11})$$

which imposes limits on π^* if inflation is to be sometimes below, in, and above the target range in REE. The higher p the less likely the Markov chain for ϵ_t is to switch state, the broader the target range possible when persistence is mild, and the narrower the range possible when persistence is high. The inequalities in (A.10) and (A.11) are more restrictive as persistence approaches its critical value $p^* = (6 + \sigma\kappa)/(6(1 + \sigma\kappa))$.

A.2 REE as a neural network

The single-layer feedforward neural network supports REE in the example nonlinear economy for:

$$\begin{aligned} a_1^* &= -\alpha\sigma\kappa\pi^*, \\ a_2^* &= \frac{(\sigma\kappa - 6p(1 + \sigma\kappa) + 6)\sigma\kappa\pi^*}{5}, \\ a_3^* &= -\frac{(\sigma\kappa - 6p(1 + \sigma\kappa) + 6)\sigma\kappa\pi^*}{5}, \\ b_{21}^* &= -\frac{6p-1}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6 + 5\alpha\sigma\kappa}, \\ b_{22}^* &= \frac{5\alpha(6p-1)}{(\sigma\kappa - 6p(1 + \sigma\kappa) + 6)(\sigma\kappa - 6p(1 + \sigma\kappa) + 6 + 5\alpha\sigma\kappa)}, \\ b_{23}^* &= -\frac{5\alpha(6p-1)}{(\sigma\kappa - 6p(1 + \sigma\kappa) + 6)(\sigma\kappa - 6p(1 + \sigma\kappa) + 6 + 5\alpha\sigma\kappa)}. \end{aligned} \quad (\text{A.12})$$

The first node is always activated and the second is activated when $a_2^* + \epsilon_t > 0$, a condition that is satisfied when $\epsilon_t > -(6p(1 + \sigma\kappa) - (6 + \sigma\kappa))\pi^*/5$. The first breakpoint in Figure 2 occurs at $\epsilon_t^* > -(6p(1 + \sigma\kappa) - (6 + \sigma\kappa))\pi^*/5 > -1$, which implies that the second node is activated when $\epsilon_t > -1$, as required in REE. Analogous reasoning shows that the third node is activated when $\epsilon_t > 1/3$, so all nodes are activated appropriately. Equations

(A.13)-(A.15) confirm that predictions from the neural network also match their REE values in Equation (A.9).

$$b_{21}^*(a_1^* + \epsilon_t) = -\frac{(6p-1)(\epsilon_t - \alpha\sigma\kappa\pi^*)}{\sigma\kappa - 6p(1+\sigma\kappa) + 6 + 5\alpha\sigma\kappa} \quad (\text{A.13})$$

$$b_{21}^*(a_1^* + \epsilon_t) + b_{22}^*(a_2^* + \epsilon_t) = -\frac{(6p-1)\epsilon_t}{\sigma\kappa - 6p(1+\sigma\kappa) + 6} \quad (\text{A.14})$$

$$b_{21}^*(a_1^* + \epsilon_t) + b_{22}^*(a_2^* + \epsilon_t) + b_{23}^*(a_3^* + \epsilon_t) = -\frac{(6p-1)(\epsilon_t + \alpha\sigma\kappa\pi^*)}{\sigma\kappa - 6p(1+\sigma\kappa) + 6 + 5\alpha\sigma\kappa} \quad (\text{A.15})$$

A.3 Prediction errors

Prediction errors depend on the differences between π_t and $E_{t-1}\pi_t$, which are functions of ϵ_t and ϵ_{t-1} in the example economy. The exogenous Markov chain for ϵ_t has six states so there are $6^2 = 36$ possible prediction errors. We collect predictions and outcomes in 6×6 matrices where the rows refer to ϵ_{t-1} and the columns to ϵ_t . Predictions $E(\Pi)$ only depend on ϵ_{t-1} in the example economy, so all the columns in (A.16) are the same.

$$E(\Pi) = \begin{pmatrix} b_{21}(a_1 - 2) \\ b_{21}(a_1 - 1) \\ b_{21}(a_1 - 1/3) + b_{22}(a_2 - 1/3) \\ b_{21}(a_1 + 1/3) + b_{22}(a_2 + 1/3) \\ b_{21}(a_1 + 1) + b_{22}(a_2 + 1) + b_{23}(a_3 + 1) \\ b_{21}(a_1 + 2) + b_{22}(a_2 + 2) + b_{23}(a_3 + 2) \end{pmatrix} \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}^T \quad (\text{A.16})$$

Outcomes depend only on ϵ_t through $E_t\pi_{t+1}$, New Keynesian Phillips curve (6), IS curve (7), and the reaction of the nominal interest rate (8). Equation (A.17) collects outcomes as a function of ϵ_{t-1} and ϵ_t in the matrix Π , where $k = 1/(1 + \alpha\sigma\kappa)$ and $E(\Pi)^T$ is a matrix of predictions conditional on ϵ_t where all rows are the same.

$$\Pi = \begin{pmatrix} k & k & 1 & 1 & k & k \\ k & k & 1 & 1 & k & k \\ k & k & 1 & 1 & k & k \\ k & k & 1 & 1 & k & k \\ k & k & 1 & 1 & k & k \\ k & k & 1 & 1 & k & k \end{pmatrix} \cdot \left[(1 + \sigma\kappa)E(\Pi)^T + \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{pmatrix} \begin{pmatrix} -\alpha\sigma\kappa\pi^* - 2 \\ -\alpha\sigma\kappa\pi^* - 1 \\ -1/3 \\ 1/3 \\ \alpha\sigma\kappa\pi^* + 1 \\ \alpha\sigma\kappa\pi^* + 2 \end{pmatrix}^T \right] \quad (\text{A.17})$$

The matrix of prediction errors is $\Pi - E(\Pi)$. When weighted by the stationary distribution of the Markov chain, the sum of the matrix elements is the unconditionally expected prediction error.

A.4 Weighting functions for updating beliefs

The 6×6 matrix of prediction errors is first weighted by the unconditional probability of each state in the stationary distribution of the Markov chain for ϵ_t . In our case, the transition probability matrix of the Markov chain is symmetric so the probability of each state is $1/36$. The second weighting is by the updating functions (1) and (2), which condition on ϵ_{t-1} and which nodes are activated when predictions are made. The weighting matrices for updating each parameter in w_1 and w_2 are in Tables A.2 and A.3. The expected update to each parameter is obtained by summing the appropriately weighted prediction errors across all states.

Parameter	Weighting matrix
a_1	$\begin{pmatrix} b_{21} & b_{21} & b_{21} & b_{21} & b_{21} & b_{21} \\ b_{21} & b_{21} & b_{21} & b_{21} & b_{21} & b_{21} \\ b_{21} & b_{21} & b_{21} & b_{21} & b_{21} & b_{21} \\ b_{21} & b_{21} & b_{21} & b_{21} & b_{21} & b_{21} \\ b_{21} & b_{21} & b_{21} & b_{21} & b_{21} & b_{21} \\ b_{21} & b_{21} & b_{21} & b_{21} & b_{21} & b_{21} \end{pmatrix}$
a_2	$\begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ b_{22} & b_{22} & b_{22} & b_{22} & b_{22} & b_{22} \\ b_{22} & b_{22} & b_{22} & b_{22} & b_{22} & b_{22} \\ b_{22} & b_{22} & b_{22} & b_{22} & b_{22} & b_{22} \\ b_{22} & b_{22} & b_{22} & b_{22} & b_{22} & b_{22} \end{pmatrix}$
a_3	$\begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ b_{23} & b_{23} & b_{23} & b_{23} & b_{23} & b_{23} \\ b_{23} & b_{23} & b_{23} & b_{23} & b_{23} & b_{23} \end{pmatrix}$

Table A.2: Weighting matrices for updating w_1

Parameter	Weighting matrix
b_{21}	$\begin{pmatrix} a_1 - 2 & a_1 - 2 & a_1 - 2 & a_1 - 2 & a_1 - 2 & a_1 - 2 \\ a_1 - 1 & a_1 - 1 & a_1 - 1 & a_1 - 1 & a_1 - 1 & a_1 - 1 \\ a_1 - 1/3 & a_1 - 1/3 & a_1 - 1/3 & a_1 - 1/3 & a_1 - 1/3 & a_1 - 1/3 \\ a_1 + 1/3 & a_1 + 1/3 & a_1 + 1/3 & a_1 + 1/3 & a_1 + 1/3 & a_1 + 1/3 \\ a_1 + 1 & a_1 + 1 & a_1 + 1 & a_1 + 1 & a_1 + 1 & a_1 + 1 \\ a_1 + 2 & a_1 + 2 & a_1 + 2 & a_1 + 2 & a_1 + 2 & a_1 + 2 \end{pmatrix}$
b_{22}	$\begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ a_2 - 1/3 & a_2 - 1/3 & a_2 - 1/3 & a_2 - 1/3 & a_2 - 1/3 & a_2 - 1/3 \\ a_2 + 1/3 & a_2 + 1/3 & a_2 + 1/3 & a_2 + 1/3 & a_2 + 1/3 & a_2 + 1/3 \\ a_2 + 1 & a_2 + 1 & a_2 + 1 & a_2 + 1 & a_2 + 1 & a_2 + 1 \\ a_2 + 2 & a_2 + 2 & a_2 + 2 & a_2 + 2 & a_2 + 2 & a_2 + 2 \end{pmatrix}$
b_{23}	$\begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ a_3 + 1 & a_3 + 1 & a_3 + 1 & a_3 + 1 & a_3 + 1 & a_3 + 1 \\ a_3 + 2 & a_3 + 2 & a_3 + 2 & a_3 + 2 & a_3 + 2 & a_3 + 2 \end{pmatrix}$

Table A.3: Weighting matrices for updating w_2

A.5 Learnability

Learnability requires that all the eigenvalues of the Jacobian of $h(\theta)$ have negative real parts when evaluated at the REE parameter values. With the determinant of the Jacobian the product of six eigenvalues, a necessary condition for agents to learn the parameters of their neural network is that the determinant is positive. We proceed by deriving necessary conditions for the determinant to be positive, then checking numerically whether they are sufficient conditions for all six eigenvalues to have positive real parts. The determinant (A.18) is obtained using the Symbolic Math Toolbox for MATLAB R2022b.

$$D = \frac{\alpha^4(\sigma\kappa)^5 \overbrace{(6p-1)^6}^{D_1 \geq 0} \overbrace{(\sigma\kappa - 6p - 4\alpha + 4\alpha p + \alpha\sigma\kappa - 6p\sigma\kappa + 4\alpha p\sigma\kappa + 6)}^{D_2 \geq 0}}{524880 \underbrace{(6p - \sigma\kappa + 6p\sigma\kappa - 6)^3}_{D_3 \geq 0} \underbrace{(\sigma\kappa - 6p + 5\alpha\sigma\kappa - 6p\sigma\kappa + 6)^3}_{D_4 > 0}} \quad (\text{A.18})$$

The determinant changes sign at two distinct values of p , the probability of the Markov chain for the disturbance term remaining in the same state in the New Keynesian Phillips curve (6). The first is when $D_3 \rightarrow \infty$, which occurs at the critical probability p^* that determines whether the disturbance term is inflationary or deflationary when inflation is in the target range. The second is at $\tilde{p}$ when $D_2 = 0$:

$$\tilde{p} = \frac{4\alpha - 6 - \sigma\kappa(1 + \alpha)}{(4\alpha - 6)(1 + \sigma\kappa)}, \quad (\text{A.19})$$

which satisfies $0 < \tilde{p} < 1$ when $\alpha > 3/2$ and $\alpha > (6 + \sigma)/(4 - \sigma)$ if $\sigma < 4$. At the interior solution, the determinant is positive for $p = 0$ and the probabilities at which the determinant changes sign are ordered by $\tilde{p} < p^*$. The determinant must therefore switch from positive to negative at $\tilde{p}$, from negative to positive at p^* , and end up positive at $p = 1$. The implication is that $p < \tilde{p}$ or $p > p^*$ are necessary conditions for learnability. Figure A.1 confirms the switching of the sign of the determinant in two numerical examples with $\pi^* = 1$. The




Figure A.1: Learnability of REE without sunspot fluctuations

light and dark grey regions in Figure A.1 support a REE with inflation sometimes below, within, and above the target range value. Values of p too low or too close to p^* (marked with dotted lines) are not allowed, because of the inequalities in (A.10) and (A.10) with $\pi^* = 1$. The numbers in bold count how many eigenvalues of the Jacobian have negative real parts in each region. Reading across at $\alpha = 3$ in the panel for $\sigma\kappa = 0.5$, there are 6 negative eigenvalues as soon as REE is supported at $p = 1/6$, but that drops to 5 when the determinant changes sign to negative at $\tilde{p} = 4/9$. The number of negative eigenvalues continues at 5 until the determinant switches again at $p^* = 13/18$, but down to 4 rather than back to 6. It follows that all six eigenvalues are only

negative in the dark grey region when $p < p^*$ and $p < \tilde{p}$. Sufficient conditions for learnability are then:¹⁶

$$p < \frac{6 + \sigma\kappa}{6(1 + \sigma\kappa)}, \quad (\text{A.20})$$

$$p < \frac{4\alpha - 6 - \sigma\kappa(1 + \alpha)}{(4\alpha - 6)(1 + \sigma\kappa)}. \quad (\text{A.21})$$

Non-interior solutions for $\tilde{p}$ exist but are not of interest. If $1 < \alpha < 2/3$ then $\tilde{p} > 1$ and $p^* < \tilde{p}$, in which case the determinant is negative at $p = 0$ and only switches to positive at p^* , violating the first condition and making it impossible to have six negative eigenvalues. If $1 < \alpha < (6 + \sigma\kappa)/(4 - \sigma\kappa)$ and $\sigma < 4$ then $\tilde{p} < 0$, the determinant is similarly negative at $p = 0$, only switching to positive at p^* and the same impossibility arises. In neither case can all the eigenvalues have negative real parts, so (A.20) and (A.21) act as necessary and sufficient conditions for learnability of the neural network parameters.

B Analytic example with sunspot fluctuations

B.1 REE

Sunspot fluctuations are accommodated in REE by expanding the state space of the example nonlinear economy. Table B.4 shows the ensuing state space, with the sunspot term taking one of two possible values $\eta_t \in (-\eta, \eta)$ when $\epsilon_t \in (-1/3, 1/3)$ and zero when $\eta_t \in (-2, -1, 1, 2)$. The idea is that sunspots will be active in REE when inflation is within the target range, but inactive when it is below or above the target range.

ϵ_t	η_t
-2	0
-1	0
-1/3	$-\eta$
-1/3	$+\eta$
1/3	$-\eta$
1/3	$+\eta$
1	0
2	0

Table B.4: Expanded state space for REE with sunspot fluctuations

The disturbance and sunspot terms follow independent exogenous Markov chains. As before, the disturbance term is unchanged with probability p and switches with probability $1 - p$. Conditional on remaining active, the sunspot term is unchanged with probability q and switches with probability $1 - q$. In terms of the expanded state space in Table B.4, the implied transition probability matrix of the Markov chain is:

$$P = \begin{pmatrix} p & (1-p)/5 & (1-p)/10 & (1-p)/10 & (1-p)/10 & (1-p)/10 & (1-p)/5 & (1-p)/5 \\ (1-p)/5 & p & (1-p)/10 & (1-p)/10 & (1-p)/10 & (1-p)/10 & (1-p)/5 & (1-p)/5 \\ (1-p)/5 & (1-p)/5 & qp & (1-q)p & q(1-p)/5 & (1-q(1-p)/5 & (1-p)/5 & (1-p)/5 \\ (1-p)/5 & (1-p)/5 & (1-q)p & qp & (1-q)(1-p)/5 & q(1-p)/5 & (1-p)/5 & (1-p)/5 \\ (1-p)/5 & (1-p)/5 & q(1-p)/5 & (1-q)(1-p)/5 & qp & (1-q)p & (1-p)/5 & (1-p)/5 \\ (1-p)/5 & (1-p)/5 & (1-q)(1-p)/5 & q(1-p)/5 & (1-q)p & qp & (1-p)/5 & (1-p)/5 \\ (1-p)/5 & (1-p)/5 & (1-p)/10 & (1-p)/10 & (1-p)/10 & (1-p)/10 & p & (1-p)/5 \\ (1-p)/5 & (1-p)/5 & (1-p)/10 & (1-p)/10 & (1-p)/10 & (1-p)/10 & (1-p)/5 & p \end{pmatrix}. \quad (\text{B.22})$$

¹⁶The second inequality relaxes in α because $d\tilde{p}/d\alpha > 0$ but tightens in $\sigma\kappa$ as $d\tilde{p}/d\sigma\kappa < 0$.

The REE with sunspot fluctuations is obtained once more by the method of undetermined coefficients. The third column of Table B.5 assigns coefficients $A, B, C_1, C_2, D_1, D_2, E, F$ to the REE value of inflation for each of the eight combinations of ϵ_t and η_t , with the remaining columns the REE prediction of inflation and the reaction of the nominal interest rate, assuming that the interest rate does not react when $\epsilon_t \in (-1/3, 1/3)$.

ϵ_t	η_t	π_t	$E_t \pi_{t+1}$	R_t
-2	0	A	$pA + ((1-p)/5)(B + E + F) + ((1-p)/10)(C_1 + C_2 + D_1 + D_2)$	$\alpha\sigma\kappa(A + \pi^*)$
-1	0	B	$pB + ((1-p)/5)(A + E + F) + ((1-p)/10)(C_1 + C_2 + D_1 + D_2)$	$\alpha\sigma\kappa(B + \pi^*)$
-1/3	$-\eta$	C_1	$p(qC_1 + (1-q)C_2) + ((1-p)/5)(A + B + qD_1 + (1-q)D_2 + E + F)$	0
-1/3	η	C_2	$p((1-q)C_1 + qC_2) + ((1-p)/5)(A + B + (1-q)D_1 + qD_2 + E + F)$	0
1/3	$-\eta$	D_1	$p(qD_1 + (1-q)D_2) + ((1-p)/5)(A + B + qC_1 + (1-q)C_2 + E + F)$	0
1/3	η	D_2	$p((1-q)D_1 + qD_2) + ((1-p)/5)(A + B + (1-q)C_1 + qC_2 + E + F)$	0
1	0	E	$pE + ((1-p)/5)(A + B + F) + ((1-p)/10)(C_1 + C_2 + D_1 + D_2)$	$\alpha\sigma\kappa(E - \pi^*)$
2	0	F	$pF + ((1-p)/5)(A + B + E) + ((1-p)/10)(C_1 + C_2 + D_1 + D_2)$	$\alpha\sigma\kappa(F - \pi^*)$

Table B.5: Expectations and nominal interest rate in REE with sunspot fluctuations

Inflation and the output gap at each ϵ_t are determined by the New Keynesian Phillips curve (6), the IS curve (7), and the REE predictions of inflation and reactions of nominal interest rates in Table B.5.

$$A = (1 + \sigma\kappa) [pA + ((1-p)/5)(B + E + F) + ((1-p)/10)(C_1 + C_2 + D_1 + D_2)] - \alpha\sigma\kappa(A + \pi^*) - 2 \quad (\text{B.23})$$

$$B = (1 + \sigma\kappa) [pB + ((1-p)/5)(A + E + F) + ((1-p)/10)(C_1 + C_2 + D_1 + D_2)] - \alpha\sigma\kappa(B + \pi^*) - 1 \quad (\text{B.24})$$

$$C_1 = (1 + \sigma\kappa) [p(qC_1 + (1-q)C_2) + ((1-p)/5)(A + B + qD_1 + (1-q)D_2 + E + F)] - 1/3 \quad (\text{B.25})$$

$$C_2 = (1 + \sigma\kappa) [p((1-q)C_1 + qC_2) + ((1-p)/5)(A + B + (1-q)D_1 + qD_2 + E + F)] - 1/3 \quad (\text{B.26})$$

$$D_1 = (1 + \sigma\kappa) [p(qD_1 + (1-q)D_2) + ((1-p)/5)(A + B + qC_1 + (1-q)C_2 + E + F)] + 1/3 \quad (\text{B.27})$$

$$D_2 = (1 + \sigma\kappa) [p((1-q)D_1 + qD_2) + ((1-p)/5)(A + B + (1-q)C_1 + qC_2 + E + F)] + 1/3 \quad (\text{B.28})$$

$$E = (1 + \sigma\kappa) [pE + ((1-p)/5)(A + B + F) + ((1-p)/10)(C_1 + C_2 + D_1 + D_2)] - \alpha\sigma\kappa(E - \pi^*) + 1 \quad (\text{B.29})$$

$$F = (1 + \sigma\kappa) [pF + ((1-p)/5)(A + B + E) + ((1-p)/10)(C_1 + C_2 + D_1 + D_2)] - \alpha\sigma\kappa(F - \pi^*) + 2 \quad (\text{B.30})$$

Equations (B.23)-(B.30) have multiple solutions that support sunspot fluctuations when

$$q = \frac{4p(1 + \sigma\kappa) + \sigma\kappa + 6}{2(4p + 1)(1 + \sigma\kappa)}, \quad (\text{B.31})$$

which is a ‘resonance frequency’ condition on the time series structure of η_t . The need for $q < 1$ imposes a further condition on p and $\sigma\kappa$, namely that $p > (4 - \sigma\kappa)/(4(1 + \sigma\kappa))$. The restricted parameter space that allows for sunspot fluctuations is in Figure B.2.




Figure B.2: Combinations of p and $\sigma\kappa$ that support sunspot fluctuations in REE

The coefficients A, B, E, F have the same values in the REE with sunspot fluctuations as the REE in Equation (A.7). The remaining coefficients C_1, C_2, D_2, D_3 are also related to their counterparts without sunspot fluctuations, except for additional dependence on the range $\pm\eta$ of fluctuations in the sunspot term.

$$\begin{aligned}
C_1 &= -\frac{5}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6} - \eta & C_2 &= -\frac{5}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6} + \eta \\
D_1 &= -\frac{5}{3(\sigma\kappa - 6p(1 + \sigma\kappa) + 6)} - \eta & D_2 &= \frac{5}{3(\sigma\kappa - 6p(1 + \sigma\kappa) + 6)} + \eta
\end{aligned}
\tag{B.32}$$

The range of the sunspot term must be small enough that inflation remains within the target range when $\epsilon_t \in (-1/3, 1/3)$. This requires $-\pi^* < C_1, C_2, D_1, D_2 < \pi^*$, which is satisfied when:

$$\eta < \frac{5}{3(\sigma\kappa - 6p(1 + \sigma\kappa) + 6)} + \pi^*
\tag{B.33}$$

The compact representation of REE with sunspot fluctuations follows the cases in Equations (A.8) and (A.9) when $\epsilon_t \in (-2, -1, 1, 2)$. If $\epsilon_t \in (-1/3, 1/3)$ then inflation and its one-period ahead prediction are:

$$\pi_t = \frac{5\epsilon_t}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6} + \eta_t,
\tag{B.34}$$

$$E_t \pi_{t+1} = -\frac{(6p - 5)\epsilon_t}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6} - \frac{\eta_t}{1 + \sigma\kappa}.
\tag{B.35}$$

B.2 REE with sunspot fluctuations as a neural network

The neural network is made compatible with sunspot fluctuations by adding the sunspot as an argument in the node that is always activated. The neural network then maps inputs $(1, \epsilon_t, \eta_t)$ to predicted inflation by:

$$E_t \pi_{t+1} = b_{21}(a_1 + b_{11}\epsilon_t + c_1\eta_t) + b_{22}\max(a_2 + b_{12}\epsilon_t, 0) + b_{23}\max(a_3 + b_{13}\epsilon_t, 0), \quad (\text{B.36})$$

where c_1 is a new parameter in the input layer. The specification supports REE when:

$$c_1^* = \frac{\sigma\kappa - 6p(1 + \sigma\kappa) + 6}{(6p - 5)(1 + \sigma\kappa)}. \quad (\text{B.37})$$

The first node is always activated, so

$$b_{21}^*(a_1^* + \epsilon_t + c_1^*\eta_t) = -\frac{(6p - 5)(\epsilon_t - \alpha\sigma\kappa\pi^*)}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6} - \frac{\eta_t}{1 + \sigma\kappa}, \quad (\text{B.38})$$

and the neural network's prediction of inflation aligns with Equation (B.35), as required.

The new parameter c_1 expands the set of parameters in the neural network that agents have to learn to $w_1 = (a_1, a_2, a_3, c_1)$ and $w_2 = (b_{21}, b_{22}, b_{23})$. The equation for updating parameter beliefs (3) now defines a 7×1 vector and the Jacobian of the associated differential equation is a 7×7 matrix. The sunspot term enters Equation (B.36) for the prediction of inflation but does not appear in the structural equations (6)-(8) for inflation, output and the nominal interest rate because it is non-fundamental.

The condition for agents to learn all the parameters of their neural network is again that the eigenvalues of the Jacobian of $h(\theta)$ have negative real parts when evaluated at the REE parameter values. This is impossible with sunspots because at least one of the eigenvalues is always zero. To see why, note that the determinant (B.39) associated with the Jacobian reduces to zero whenever the 'resonance frequency' condition (B.31) needed to support sunspot fluctuations in REE is satisfied.

$$D \propto (2(4p + 1)(1 + \sigma\kappa)q - (4p(1 + \sigma\kappa) + \sigma + 6)) \quad (\text{B.39})$$

With the determinant of the Jacobian equal to zero, at least one eigenvalue always has a non-negative real part and agents cannot learn all the parameters of their neural network.

C Analytic example with inflation always outside target range

C.1 REE

REE exist in which inflation is outside its target range for any value of the disturbance term. We classify these equilibria using the notation $(\underline{n}, n, \bar{n})$ for the number of states with inflation below, within and above the target range. Our interest in this subsection of the appendix is in REE with $(\underline{n}, 0, 6 - \underline{n})$ for $\underline{n} = 0 \dots 6$.¹⁷ Candidate REE are identified by assuming that inflation is below or above the target range in the required number of states and applying the method of undetermined coefficients. There is some commonality in the

¹⁷The main analysis of the example nonlinear economy concerns REE of type (2,2,2).

specification of candidate REEs across the cases, so Table C.6 eases presentation by defining:

$$\phi = \frac{5}{\sigma\kappa - 6p(1 + \sigma\kappa) + 6 + 5\alpha\sigma\kappa} \quad (\text{C.40})$$

$$\phi^0 = \frac{\alpha}{\alpha - 1} \quad (\text{C.41})$$

$$\phi^1 = -\frac{\phi^0\phi}{5}(4p + \sigma\kappa - 5\alpha\sigma\kappa + 4p\sigma\kappa - 4) \quad (\text{C.42})$$

$$\phi^2 = -\frac{\phi^0\phi}{5}(2p - 2 - 5\alpha\sigma\kappa + 3\sigma\kappa + 2p\sigma\kappa) \quad (\text{C.43})$$

$$\phi^3 = \phi\alpha\sigma\kappa \quad (\text{C.44})$$

$$\Delta = 2\phi\alpha\sigma\kappa \quad (\text{C.45})$$

Coefficient	(0, 0, 6)	(1, 0, 5)	(2, 0, 4)	(3, 0, 3)
A	$-2\phi + \phi^0\pi^*$	$-2\phi + (\phi^1 - \Delta)\pi^*$	$-2\phi + (\phi^2 - \Delta)\pi^*$	$-2\phi + (\phi^3 - \Delta)\pi^*$
B	$-\phi + \phi^0\pi^*$	$-\phi + \phi^1\pi^*$	$-\phi + (\phi^2 - \Delta)\pi^*$	$-\phi + (\phi^3 - \Delta)\pi^*$
C	$-\phi/3 + \phi^0\pi^*$	$-\phi/3 + \phi^1\pi^*$	$-\phi/3 + \phi^2\pi^*$	$-\phi/3 + (\phi^3 - \Delta)\pi^*$
D	$\phi/3 + \phi^0\pi^*$	$\phi/3 + \phi^1\pi^*$	$\phi/3 + \phi^2\pi^*$	$\phi/3 + \phi^3\pi^*$
E	$\phi + \phi^0\pi^*$	$\phi + \phi^1\pi^*$	$\phi + \phi^2\pi^*$	$\phi + \phi^3\pi^*$
F	$2\phi + \phi^0\pi^*$	$2\phi + \phi^1\pi^*$	$2\phi + \phi^2\pi^*$	$2\phi + \phi^3\pi^*$

Table C.6: Candidate REEs with inflation always outside target range

Inflation reacts to the disturbance term through ϕ in the same way in all candidate REE, with the same intensity seen outside the inflation target range in (2,2,2) equilibria without and with sunspot fluctuations in Equation (9). The constant term varies depending on the number of states below or above the target range, but it always jumps by the same amount Δ when inflation is assumed to go from below to above. Figure C.3 shows a numerical example, with $p = 3/4$, $\alpha = 5$, $\pi^* = 1$ and $\sigma\kappa = 1$.¹⁸




Figure C.3: Numerical example candidate REEs with inflation always outside target range

¹⁸Candidate REE for (4,0,2), (5,0,1) and (6,0,0) are obtained by symmetry.

It remains to verify whether the candidate REEs are consistent with the number of states where inflation is assumed to be below and above the target range. A first check is that the jump in the constant term has to be greater than the width of the range, i.e. $\Delta > 2\pi^*$. From Equation (C.40), this requires $p > p^*$ so there is no alternative (1,0,5), (2,0,4) or (3,0,3) REE in the configuration with mild persistence in Figure 2a. A second check is whether the jumps in the constant term allow inflation to straddle the target range appropriately. The (3,0,3) candidate passes trivially as the jump in inflation is centred on zero, and there is always an alternative REE in the configuration with high persistence in Figure 2b. The (6,0,0) candidate does not straddle the target range but requires $-2\phi + \phi^0\pi^* > \pi^*$ to be always above it. The condition depends on parameter values but is never satisfied for $p < p^*$, so there is also no alternative (0,0,6) REE in Figure 2a. Existence of REE for (1,0,5) and (2,0,4) also depends on parameter values, although if an REE exists for $(\underline{n}, 0, 6 - \underline{n})$ and $\underline{n} \in (0, 1, 2)$ then it also exists for $(\underline{n} + 1, 0, 5 - \underline{n})$. The parameterisation in Figure C.3 permits (2,0,4), (3,0,3) and (4,0,2) as REE with inflation always outside the target range.

C.2 REE as a neural network

The compact representation of the (3,0,3) REE:

$$\pi_t = \begin{cases} \phi\epsilon_t - \phi\alpha\sigma\kappa\pi^* & \text{if } \epsilon_t \in (-2, -1, -1/3), \\ \phi\epsilon_t + \phi\alpha\sigma\kappa\pi^* & \text{if } \epsilon_t \in (1/3, 1, 2), \end{cases}$$

$$E_t\pi_{t+1} = \begin{cases} \left(\frac{6p-1}{5}\right)(\phi\epsilon_t - \phi\alpha\sigma\kappa\pi^*) & \text{if } \epsilon_t \in (-2, -1, -1/3), \\ \left(\frac{6p-1}{5}\right)(\phi\epsilon_t + \phi\alpha\sigma\kappa\pi^*) & \text{if } \epsilon_t \in (1/3, 1, 2). \end{cases}$$

is compatible with multiple parameterisations of the neural network because inflation is never in the target range to identify precisely how equilibrium inflation jumps across the target range. Without loss of generality, we therefore normalise the jump to occur linearly between $\epsilon_t = -1/6$ and $\epsilon_t = 1/6$, and note that agents still have to learn the parameters of their neural network when inflation is below and above the target range. The parameter values consistent with the (3,0,3) normalised REE are:

$$\begin{aligned} a_1^* &= -\alpha\sigma\kappa\pi^*, & a_2^* &= \frac{1}{6}, & a_3^* &= -\frac{1}{6}, \\ b_{21}^* &= \frac{(6p-1)\phi}{5}, & b_{22}^* &= \frac{6(6p-1)\phi\alpha\sigma\kappa\phi^*}{5}, & b_{23}^* &= -\frac{6(6p-1)\phi\alpha\sigma\kappa\phi^*}{5}, \end{aligned} \tag{C.46}$$

which are consistent with REE because:

$$b_{21}^*(a_1^* + \epsilon_t) = \left(\frac{6p-1}{5}\right)(\phi\epsilon_t - \phi\alpha\sigma\kappa\pi^*), \tag{C.47}$$

$$b_{21}^*(a_1^* + \epsilon_t) + b_{22}^*(a_2^* + \epsilon_t) = \left(\frac{6p-1}{5}\right)\left(\phi + \frac{6\phi\alpha\sigma\kappa\pi^*}{5}\right), \tag{C.48}$$

$$b_{21}^*(a_1^* + \epsilon_t) + b_{22}^*(a_2^* + \epsilon_t) + b_{23}^*(a_3^* + \epsilon_t) = \left(\frac{6p-1}{5}\right)(\phi\epsilon_t + \phi\alpha\sigma\kappa\pi^*). \tag{C.49}$$

The first node in the neural network is always active, the second active when $\epsilon_t > -1/6$, and the third when $\epsilon_t > 1/6$. In equilibrium, either just the first node is active because $\epsilon_t < -1/6$ or all three nodes are active because $\epsilon_t > 1/6$. There are no observations of inflation in the target range that separately identify the activation of the second and third nodes, hence the need to normalise the neural network.

C.3 Learnability

Agents have four parameters to learn in their neural network, a_1 and b_{21} from the first node and a_3 and b_{23} from the third node. The ordinary differential equation $h(\theta)$ associated with updating parameter estimates is hence a 4×1 vector with a 4×4 Jacobian. Learnability requires all four eigenvalues of the Jacobian to have negative real parts when evaluated at the REE parameter values, a condition only satisfied if its determinant is positive. Symbolic Math Toolbox for MATLAB R2022b provides the determinant:

$$D = -\frac{19\alpha^2(\pi^*)^2(\sigma\kappa)^3(6p-1)^4(\alpha-1)}{20250(\alpha\sigma\kappa+1)^4(\sigma\kappa-6p+5\alpha\sigma\kappa-6p\sigma\kappa+6)^2} > 0 \quad (\text{C.50})$$

The expression is always positive so either all four eigenvalues are negative or two are negative and two are positive. The latter is ruled out by considering the characteristic polynomial of the Jacobian, $\lambda^4 + d_3\lambda^3 + d_2\lambda^2 + d_1\lambda + d_0$. The coefficient on λ^4 is positive by construction and d_0 is positive because it equals the determinant. Further calculations using the Symbolic Math Toolbox show that d_3 , d_2 and d_1 are positive for any values of the structural parameters ($0 < p < 1, \alpha > 1, \pi^* > 0, \sigma\kappa > 0$). The coefficients of the characteristic polynomial thus do not change sign with powers of λ , and by Descartes' rule of signs the Jacobian has no positive roots. It follows that all eigenvalues are negative, the REE outside the target range is always learnable, and there is no need to check learnability at different values of the structural parameters.

D Numerical example analytics

D.1 Deterministic steady states

The steady-state New Keynesian Phillips curve (16) and combined IS (17) and policy rule (18) are:

$$\pi = \beta\pi + \kappa y, \quad (\text{D.51})$$

$$y = \begin{cases} \eta y - \sigma(\phi_\pi + \alpha - 1)\pi - \sigma\alpha\pi^* & \text{if } \pi < -\pi^*, \\ \eta y - \sigma(\phi_\pi - 1)\pi & \text{if } -\pi^* \leq \pi \leq \pi^*, \\ \eta y - \sigma(\phi_\pi + \alpha - 1)\pi + \sigma\alpha\pi^* & \text{if } \pi > \pi^*. \end{cases} \quad (\text{D.52})$$

There is always a deterministic steady state at (0,0), and multiple deterministic steady states if the steady-state combined IS and policy rule is flatter in the inflation target range than the steady-state Phillips curve. Multiplicity occurs when inertia in the IS curve is high, when the coefficient on lagged inflation satisfies:

$$\eta > 1 - \frac{\sigma\kappa(1 - \phi_\pi)}{1 - \beta}. \quad (\text{D.53})$$

D.2 Local determinacy properties of deterministic steady states

The local determinacy properties in the neighbourhood of a deterministic steady state are governed by the eigenvalues of the characteristic polynomial of the linearised model. In the neighbourhood of the (0,0) deterministic steady state, the linearised model has the state space form:

$$E_t \begin{pmatrix} \beta & \kappa \\ \sigma & -1 \end{pmatrix} \begin{pmatrix} \pi_{t+1} \\ y_t \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ \sigma\phi_\pi & -\eta \end{pmatrix} \begin{pmatrix} \pi_t \\ y_{t-1} \end{pmatrix} + \begin{pmatrix} \epsilon_{\pi,t} \\ \epsilon_{y,t} \end{pmatrix}, \quad (\text{D.54})$$

and the characteristic polynomial is:

$$(\beta + \sigma\kappa)\lambda_1^2 - (1 + \sigma\kappa\phi_\pi + \beta\eta)\lambda_1 + \eta = 0. \quad (\text{D.55})$$

Following Sargent (1987, p. 187-189), the (0,0) steady state is indeterminate if and only if:

$$\eta > 1 - \frac{\sigma\kappa(1 - \phi_\pi)}{1 - \beta}, \quad (\text{D.56})$$

which is the same condition as that for the multiplicity of deterministic steady states. If inertia is low then the only deterministic steady state is (0,0) and it is locally determinate. If inertia is high then there are multiple deterministic steady states and the deterministic steady state at (0,0) is locally indeterminate. In the latter case, the deterministic steady states other than (0,0) are always locally determinate.

E Example with Zero Lower Bound

E.1 Equilibrium conditions

The Lagrangian of the household-producer (E.57) takes (R_t, g_t, A_t, ξ_t) as given:

$$\mathcal{L} = E_{0,i} \sum_{t=0}^{\infty} \beta^t \left\{ \begin{aligned} & \log(c_{t,i} + \xi_t g_t - \lambda c_{t-1}) + \chi \log\left(\frac{M_{t-1,i}}{P_t}\right) - (1+\epsilon)^{-1} h_{t,i}^{1+\epsilon} - \Phi\left(\frac{P_{t,i}}{P_{t-1,i}}\right) \\ & + \mu_{t,i} \left(\frac{m_{t-1,i}}{\pi_t} + R_{t-1} \frac{b_{t-1,i}}{\pi_t} + \frac{P_{t,i}}{P_t} y_{t,i} - c_{t,i} - m_{t,i} - b_{t,i} + \Upsilon_{t,i} \right) \\ & + \Lambda_{t,i} c_{t,i} \\ & + \gamma_{t,i}^0 (A_t h_{t,i}^\alpha - y_{t,i}) \\ & + \gamma_{t,i}^1 \left(\left(\frac{y_{t,i}}{y_t} \right)^{-1/\nu} P_t - P_{t,i} \right) \end{aligned} \right\}. \quad (\text{E.57})$$

The first order conditions with respect to $(c_{t,i}, b_{t,i}, h_{t,i}, P_{t,i}, y_{t,i})$ are:

$$0 = \frac{1}{c_{t,i} + \xi_t g_t - \lambda c_{t-1}} + \Lambda_{t,i} - \mu_{t,i} \quad (\text{E.58})$$

$$0 = -\mu_{t,i} + \beta R_t E_{0,i} \left(\frac{\mu_{t+1,i}}{\pi_{t+1}} \right) \quad (\text{E.59})$$

$$0 = -h_{t,i}^\epsilon + \gamma_{t,i}^0 \alpha A_t h_{t,i}^{\alpha-1} \quad (\text{E.60})$$

$$0 = -\Phi' \left(\frac{P_{t,i}}{P_{t-1,i}} \right) \frac{1}{P_{t-1,i}} + \mu_{t,i} \frac{y_{t,i}}{P_t} - \gamma_{t,i}^1 + \beta E_{0,i} \left(\Phi' \left(\frac{P_{t+1,i}}{P_{t,i}} \right) \frac{P_{t+1,i}}{P_{t,i}^2} \right) \quad (\text{E.61})$$

$$0 = \mu_{t,i} \frac{P_{t,i}}{P_t} - \gamma_{t,i}^0 - \gamma_{t,i}^1 \frac{1}{\nu} \left(\frac{y_{t,i}}{y_t} \right)^{-1/\nu-1} \frac{P_t}{y_t} \quad (\text{E.62})$$

Given all household-producers are identical, we drop i subscripts from the first order conditions:

$$0 = \frac{1}{c_t + \xi_t g_t - \lambda c_{t-1}} + \Lambda_t - \mu_t \quad (\text{E.63})$$

$$0 = -\mu_t + \beta R_t E_t \left(\frac{\mu_{t+1}}{\pi_{t+1}} \right) \quad (\text{E.64})$$

$$0 = -h_t^\epsilon + \gamma_t^0 \alpha A_t h_t^{\alpha-1} \quad (\text{E.65})$$

$$0 = -\Phi'(\pi_t) \pi_t + \mu_t y_t - \gamma_t^1 P_t + \beta E_t (\Phi'(\pi_{t+1}) \pi_{t+1}) \quad (\text{E.66})$$

$$0 = \mu_t - \gamma_t^0 - \frac{\gamma_t^1 P_t}{\nu y_t} \quad (\text{E.67})$$

$\Lambda_t = 0$ when the non-negativity constraint on consumption is non-binding. Combining equations (E.63) and (E.64) then gives the unconstrained Euler equation, which is Equation (30) in the paper:

$$\frac{1}{c_t + \xi_t g_t - \lambda c_{t-1}} = \beta R_t E_t \left(\frac{\pi_{t+1}^{-1}}{c_{t+1} + \xi_{t+1} g_{t+1} - \lambda c_t} \right) \quad (\text{E.68})$$

Irrespective of whether the non-negativity constraint on consumption binds, Equations (E.65), (E.66), (E.67) and market clearing $y_t = c_t + g_t$ define the New Keynesian Phillips Curve, Equation (32) in the paper:

$$\Phi'(\pi_t) \pi_t = \frac{\nu}{\alpha} \left(\frac{c_t + g_t}{A_t} \right)^{\frac{\epsilon+1}{\alpha}} + \frac{1-\nu}{c_t + \xi_t g_t - \lambda c_{t-1}} (c_t + g_t) + \beta E_t \Phi'(\pi_{t+1}) \pi_{t+1} \quad (\text{E.69})$$

The exogenous shocks (A_t, ξ_t) follow AR(1) processes with Gaussian innovations:

$$A_t = A + \rho_A A_{t-1} + \epsilon_{A,t} \quad \text{where} \quad \epsilon_{A,t} \sim \mathcal{N}(0, \sigma_A) \quad (\text{E.70})$$

$$\xi_t = \xi + \rho_\xi \xi_{t-1} + \epsilon_{\xi,t} \quad \text{where} \quad \epsilon_{\xi,t} \sim \mathcal{N}(0, \sigma_\xi) \quad (\text{E.71})$$

E.2 Deterministic steady states

For the unconstrained Euler equation (E.68) to hold in a steady state, it must be that $\pi = \beta R$. Since $R \geq 1$ is subject to the zero lower bound, steady-state consumption will only be unconstrained when $\pi \geq \beta$, in which case total differentiation of the nominal interest rate rule yields:

$$\frac{d\pi}{dy} = \frac{\phi^y}{1 - \phi^\pi} \frac{\pi}{y} \leq 0, \quad (\text{E.72})$$

and the steady-state Euler equation is downward-sloping. When $\pi = \beta$, the steady-state nominal interest rate is constrained at $R = 1$ and steady-state output can take any value between a lower bound $\underline{y}$ supported by government spending and an upper bound $\bar{y}$ at which R is no longer constrained. The bounds satisfy:

$$\underline{y} = \frac{\bar{g}}{1 + e^{k(\beta - \pi^*)}}, \quad (\text{E.73})$$

$$R^* \left(\frac{\beta}{\pi^*} \right)^{\phi_\pi} \left(\frac{\bar{y}}{y^*} \right)^{\phi_y} = 1. \quad (\text{E.74})$$

The non-negativity constraint on consumption binds whenever $\pi < \beta$. In this case, steady-state consumption is zero and steady-state output is entirely determined by government spending. By total differentiation:

$$\frac{d\pi}{dy} = - \frac{(1 + e^{k(\pi - \pi^*)})^2}{\bar{g} k e^{k(\pi - \pi^*)}} \leq 0, \quad (\text{E.75})$$

and the steady-state Euler equation is downward-sloping. The steady-state Phillips Curve:

$$(1 - \beta)\Phi'(\pi)\pi = \frac{\nu}{\alpha} \left(\frac{y}{A} \right)^{\frac{\epsilon+1}{\alpha}} + \frac{1 - \nu}{(1 - \lambda)y + (\xi - 1 + \lambda) \frac{\bar{g}}{1 + e^{k(\pi - \pi^*)}}} y, \quad (\text{E.76})$$

is upward-sloping for an appropriate specification of the pricing friction $\Phi(\cdot)$ (Evans et al., 2022). Deterministic steady states occur at the intersections of the steady state versions of the Euler equation and Phillips Curve.

E.3 Parameter values

Following Evans et al. (2022), the pricing friction is a Linex function:

$$\Phi \left(\frac{P_{t,i}}{P_{t-1,i}} \right) \equiv \frac{\phi}{\psi^2} \left[\exp \left(-\psi \left(\left(\frac{P_{t,i}}{P_{t-1,i}} \right) - \pi^* \right) \right) + \psi \left(\left(\frac{P_{t,i}}{P_{t-1,i}} \right) - \pi^* \right) - 1 \right] \quad (\text{E.77})$$

β	π^*	α	$\bar{\xi}$	$\bar{A}$	ν	$\bar{g}$	ϕ	ψ
0.96	1.02	0.35	1	1	2	0.8	75	5
ϵ	ϕ^π	ϕ^y	λ	k	ρ_A	ρ_ξ	σ_A	σ_ξ
0	2	1	0.95	2	0.5	0.5	0.01	0.1

Table E.7: Parameter values for the ZLB example

The household-producer is unconstrained at the target steady state, so $R^* = \pi^*/\beta = 1.0625$. As in Evans et al. (2022), $\Phi'(\pi)\pi = 0$ at π^* and the steady-state Phillips Curve pins down $y^* = 0.631$ compatible with π^* .

References

- George-Marios Angeletos, Fabrice Collard, and Harris Dellas. Public Debt as Private Liquidity: Optimal Policy. NBER Working Papers 22794, National Bureau of Economic Research, 2021.
- S Borağan Aruoba, Pablo Cuba-Borda, and Frank Schorfheide. Macroeconomic Dynamics Near the ZLB: A Tale of Two Countries. *Review of Economic Studies*, 85(1), 2018.
- S Borağan Aruoba, Pablo Cuba-Borda, and Frank Schorfheide. Macroeconomic Dynamics Near the ZLB: A Tale of Two Countries. *Review of Economic Studies*, 85(1):87–118, 2017.
- Jess Benhabib, Stephanie Schmitt-Grohé, and Martin Uribe. The Perils of Taylor Rules. *Journal of Economic Theory*, 96(1-2):40–69, 2001.
- Pierpaolo Benigno and Gauti B. Eggertsson. It’s Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve. NBER Working Papers 31197, National Bureau of Economic Research, 2023.
- Francesco Bianchi, Leonardo Melosi, and Matthias Rottner. Hitting the elusive inflation target. *Journal of Monetary Economics*, 124:107–122, 2021.
- Olivier Jean Blanchard and Charles M Kahn. The Solution of Linear Difference Models under Rational Expectations. *Econometrica: Journal of the Econometric Society*, pages 1305–1311, 1980.
- R Anton Braun and Lena Mareen Körber. New Keynesian dynamics in a low interest rate environment. *Journal of Economic Dynamics and Control*, 35(12):2213–2227, 2011.
- James Bullard. Learning Equilibria. *Journal of Economic Theory*, 64(2):468–485, 1994.
- James Bullard. Seven Faces of “The Peril”. *Federal Reserve Bank of St. Louis Review*, 92(September/October), 2010.
- Guido Cazzavillan, Teresa Lloyd-Braga, and Patrick A Pintus. Multiple Steady States and Endogenous Fluctuations with Increasing Returns to Scale in Production. *Journal of Economic Theory*, 80(1):60–107, 1998.
- Lawrence Christiano, Martin S Eichenbaum, and Benjamin K Johannsen. Does the New Keynesian Model Have a Uniqueness Problem? NBER Working Papers 24612, National Bureau of Economic Research, 2018.
- George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.
- Wouter J Den Haan and Albert Marcet. Accuracy in Simulations. *Review of Economic Studies*, 61(1):3–17, 1994.
- Jan Eeckhout and Ilse Lindenlaub. Unemployment Cycles. *American Economic Journal: Macroeconomics*, 11(4):175–234, 2019.
- Gauti B Eggertsson and Michael Woodford. The Zero Bound on Interest Rates and Optimal Monetary Policy. *Brookings Papers on Economic Activity*, 2003(1):139–233, 2003.
- Gauti B Eggertsson, Neil R Mehrotra, and Jacob A Robbins. A Model of Secular Stagnation: Theory and Quantitative Evaluation. *American Economic Journal: Macroeconomics*, 11(1):1–48, 2019.
- Martin Ellison and Joseph Pearlman. Saddlepath Learning. *Journal of Economic Theory*, 146(4):1500–1519, 2011.
- George W Evans and Seppo Honkapohja. *Learning and Expectations in Macroeconomics*. Princeton University Press, 2001.
- George W Evans, Eran Guse, and Seppo Honkapohja. Liquidity traps, learning and stagnation. *European Economic Review*, 52(8):1438–1463, 2008.
- George W Evans, Seppo Honkapohja, and Kaushik Mitra. Expectations, Stagnation and Fiscal Policy: a Nonlinear Analysis. *International Economic Review*, 63(3):1397–1425, 2022.

- Roger EA Farmer, Vadim Khramov, and Giovanni Nicolò. Solving and Estimating Indeterminate DSGE models. *Journal of Economic Dynamics and Control*, 54:17–36, 2015.
- Kristin Forbes, Joseph Gagnon, and Christopher G. Collins. Low Inflation Bends the Phillips Curve around the World. NBER Working Papers 29323, National Bureau of Economic Research, 2021.
- Jordi Galí. Product diversity, endogenous markups, and development traps. *Journal of Monetary Economics*, 36(1):39–63, 1995.
- Alfred Greiner and Anton Bondarev. Optimal R&D investment with learning-by-doing: Multiple steady states and thresholds. *Optimal Control Applications and Methods*, 38(6):956–962, 2017.
- Luca Guerrieri and Matteo Iacoviello. Occbin: A toolkit for solving dynamic models with occasionally binding constraints easily. *Journal of Monetary Economics*, 70:22–38, 2015.
- Pablo A Guerrón-Quintana and James M Nason. Bayesian Estimation of DSGE Models. *Handbook of Research Methods and Applications in Empirical Macroeconomics*, pages 486–512, 2013.
- James D Hamilton. A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica*, 57(2):357–384, 1989.
- Cars Hommes and Gerhard Sorger. Consistent Expectations Equilibria. *Macroeconomic Dynamics*, 2(3):287–321, 1998.
- Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2):251–257, 1991.
- Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989.
- Greg Kaplan and Guido Menzio. Shopping Externalities and Self-Fulfilling Unemployment Fluctuations. *Journal of Political Economy*, 124(3):771–825, 2016.
- Hanno Kase, Leonardo Melosi, and Matthias Rottner. Estimating Nonlinear Heterogeneous Agents Models with Neural Networks. Working Papers 2022-26, Federal Reserve Bank of Chicago, 2022.
- Paul Krugman. History Versus Expectations. *Quarterly Journal of Economics*, 106(2):651–667, 1991.
- Hervé Le Bihan, Magali Marx, and Julien Matheron. Inflation tolerance ranges in the New Keynesian model. *European Economic Review*, 153, 2023.
- Thomas A Lubik and Frank Schorfheide. Testing for Indeterminacy: An Application to U.S. Monetary Policy. *American Economic Review*, 94(1):190–217, 2004.
- Karl-Göran Mäler, Anastasios Xepapadeas, and Aart De Zeeuw. The Economics of Shallow Lakes. *Environmental and Resource Economics*, 26(4):603–624, 2003.
- Lilia Maliar, Serguei Maliar, and Pablo Winant. Deep learning for solving dynamic economic models. *Journal of Monetary Economics*, 122:76–101, 2021.
- Carolina Manzano and Xavier Vives. Public and private learning from prices, strategic substitutability and complementarity, and equilibrium multiplicity. *Journal of Mathematical Economics*, 47(3):346–369, 2011.
- Taisuke Nakata and Sebastian Schmidt. Conservatism and liquidity traps. *Journal of Monetary Economics*, 104:37–47, 2019.
- Alexander W Richter and Nathaniel A Throckmorton. The zero lower bound: frequency, duration, and numerical convergence. *BE Journal of Macroeconomics*, 15(1):157–182, 2015.
- Thomas J Sargent. *Macroeconomic Theory*. Academic Press, 2nd edition, 1987.
- Christopher A Sims. Solving Linear Rational Expectations Models. *Computational Economics*, 20(1):1–20, 2002.