

The Unattractiveness of Indeterminate Dynamic Equilibria*

Julian Ashwin Paul Beaudry Martin Ellison[†]
London Business School Bank of Canada University of Oxford and CEPR

December 21, 2022

Abstract

Macroeconomic forces that generate multiple equilibria often support locally-indeterminate dynamic equilibria in which a continuum of perfect foresight paths converge towards the same steady state. The set of rational expectations equilibria (REE) in such environments can be very large, although the relevance of many of them has been questioned on the basis that they may not be learnable. In this paper we document the existence of a learnable REE in such situations. However, we show that the dynamics of this learnable REE do not resemble perturbations around any of the convergent perfect foresight paths. Instead, the learnable REE treats the locally-indeterminate steady state as unstable, in contrast to it resembling a stable attractor under perfect foresight.

Keywords: indeterminacy, machine learning, multiple equilibria, neural networks

*The authors gratefully acknowledge discussions with Guido Ascari, Roger Farmer, Michael McMahon, Galo Nuño, Anna Rogantini Picco, 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.

[†]Corresponding author: martin.ellison@economics.ox.ac.uk

1 Introduction

Macroeconomic forces often generate multiple equilibria and indeterminacy under rational expectations. Well-known examples include increasing returns to scale in a production function (Benhabib and Farmer, 1994; Cazzavillan et al., 1998); a passive monetary policy rule (Clarida et al., 2000; Lubik and Schorfheide, 2004); an Effective Lower Bound on monetary policy (Benhabib et al., 2001; Christiano et al., 2018; Eggertsson et al., 2019); externalities in a search and matching framework in the goods (Huo and Ríos-Rull, 2013; Kaplan and Menzio, 2016), labour (Eeckhout and Lindenlaub, 2019) or inter-firm market (Fernandez-Villaverde et al., 2020); endogenous markups generating non-convex marginal revenue product of capital (Gali, 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 (Måler et al., 2003); positive externalities in production (Krugman, 1991); feedback between government debt and interest rate through liquidity constraint (Angeletos et al., 2021).

Many of these papers recognise that the non-convexity or complementarity driving indeterminacy is likely to be only locally powerful, in which case there will be multiple steady states with regions of local indeterminacy around some of the steady states. The steady states around which there is indeterminacy feature a continuum of converging stable paths, all of which are consistent with rationality and so qualify as rational expectations equilibria (REE). The set of these REE is typically very large, especially if the environment has stochastic elements. To render the multiplicity manageable, a common approach in many of the papers above is to focus on REE that are perturbations of perfect foresight solutions of the model, and then invoke sunspots to select between the many candidate stationary paths in locally-indeterminate regions.

In this paper, we argue against the focus on REE that are perturbations of perfect foresight paths. Instead, we propose that the multiplicity of REE should be resolved by appealing to learnability. There are strong precedents for using a learnability criterion to discriminate between multiple candidate REE in indeterminate systems (Evans and Honkapohja, 2001; Woodford, 2003), although the learning literature has found it difficult to derive general conditions for learnability in non-linear models with multiple steady states.¹ In particular, in environments where there are a continuum of perfect foresight paths that converge to a steady state, the literature has not yet established whether there is a learnable REE near that steady state and, if there is one, what it resembles.

Given the tensions between perfect foresight paths, rational expectations, and learnability, we set up a prototypical environment and ask whether it possesses a REE that is learnable. The environment we build upon involves monetary policy in a sticky price setting, but many other interpretations could be given. In our illustrative setup, the nominal interest rate is assumed to react less than one-for-one to inflation when it is close to target, breaking the Taylor Principle and generating a region of local indeterminacy around a steady state centred on the inflation target, as in Clarida et al. (2000). The reaction of the nominal interest rate is assumed

¹See Section 12.6 in Evans and Honkapohja (2001) and Section 2.3 in Woodford (2003). Sunspot equilibria have been shown to be learnable under restrictive conditions (Woodford, 1990; McCallum, 2007; Ellison and Pearlman, 2011), although Arifovic et al. (2013) conclude that “In the recursive learning literature, it has generally been difficult to obtain expectational stability of sunspot equilibria.”

to strengthen as inflation moves away from its target, eventually satisfying the Taylor Principle and supporting two further steady states around which equilibrium dynamics are locally determinate. Since our exploration is numerical, our notion of an REE corresponds to that proposed by Den Haan and Marcet (1994), whereby forecast errors in rational expectations equilibrium should be unpredictable. In searching for a learnable REE, we allow agents to be quite sophisticated in that we allow them to exploit neural networks to form expectations.

We find a learnable REE exists in our prototypical environment. However, its properties do not resemble simple perturbations around convergent perfect foresight paths. Instead of multiple perfect-foresight paths converging to the steady state of the locally-indeterminate region, there is a unique REE path that leads away from the locally-indeterminate region towards one of the two steady states that are surrounded by locally-determinate dynamics. If one accepts learnability as a good equilibrium selection device, our results suggest that a steady state with a continuum of convergent perfect foresight paths around it (a sink steady state) should be treated as an unstable equilibrium. Where shocks may land one near this steady state, equilibrium behaviour pushes away from it, causing the system to spend very little time near the steady state of interest.

The use of neural network learning to derive our results raises the question of whether there would be convergence to the same equilibrium behaviour if agents were less sophisticated. Our answer is a qualified yes. To be more precise, we show that if agents were less sophisticated and used adaptive learning then their equilibrium behaviour converges toward the same type of REE configuration. Their forecast errors would exhibit some slight predictability because they are less sophisticated, meaning that their behaviour would not exactly converge to the REE we identify. It would nevertheless be very close, with equilibrium dynamics having the same key properties. Given these results, we believe our approach identifies the most probable REE for the environment considered.

Simulations of the learnable REE from the prototypical model resemble those from a simple Markov-switching process, and standard econometric tests on simulated data identify two distinct regimes even though the learnable REE is unique. The danger of improper econometric inference extends to tests for indeterminacy based on estimating log-linearised dynamics around a steady state (Lubik and Schorfheide, 2004). An econometrician running these tests could incorrectly infer that economic dynamics are characterised by convergent behaviour perturbed by sunspot shocks, when in truth they are governed by the learnable REE where behaviour is divergent and there are no sunspot shocks.

The paper is structured as follows. In Section 2 we discuss the related literature to set the scene for the prototypical setup we consider in Section 3. We focus on an environment that has three deterministic steady states, the central one being a locally indeterminate sink under rational expectations and the outer two being locally determinate saddles. Neural network learning and our concept of a learnable REE are introduced in Section 4. Our main results are in Section 5, which describes how the locally-indeterminate sink steady state becomes a source in the learnable REE identified by neural network learning. The robustness of this result is demonstrated in Section 6, which confirms similar equilibrium dynamics if agents form expectations using less sophisticated adaptive learning techniques. Section 7 addresses the question of whether an econometrician

testing for local indeterminacy could mistakenly infer that data is generated by sunspot shocks. An application of our findings to an environment with an effective lower bound on nominal interest rates is in Section 8, before we conclude with Section 9.

2 Related literature

The literature on learning in macroeconomics goes back at least as far as Lucas (1986), who suggested that stability under learning could be used as a criterion to select between REE in an overlapping generations model of fiat money. The most influential contribution is Evans and Honkapohja (2001), which provides a rich and powerful analytic framework for analysing learning. Their concept of E-stability is used by McCallum (2007) and Ellison and Pearlman (2011) to verify that REE around locally-indeterminate steady states are not generally learnable, except under restrictive conditions. For example, agents can only learn the location of the steady state around which equilibria are locally-indeterminate if there are strong negative feedbacks from expectations to outcomes, such as in the model of Evans and McGough (2005). Learnable sunspot equilibria have been found to exist in specific circumstances (Woodford, 1990; Arifovic et al., 2013) but the general expectation stability of sunspot equilibria has proved elusive (Evans and Honkapohja, 2001; Woodford, 2003). If learnability is interpreted as satisfying the stricter criterion of iterative e-stability then locally-indeterminate REE are never learnable (Woodford, 2003).²

A comparatively small literature tackles learning in non-linear models, where exact E-stability analysis in the spirit of Evans and Honkapohja (2001) is not possible because REE in a non-linear model with continuous endogenous state variables cannot be expressed in a finite number of parameters.³ Some papers deal with non-linearity by modelling linear learning in a non-linear model (Hommes and Sorger, 1998; Bullard, 1994), although expectations in these equilibria will never be consistent with rationality. Berardi and Duffy (2015) use a PEA approach to approximate non-linear expected values around a steady state, showing reasonable convergence results but not evidencing the accuracy of the learning equilibrium.

Although a well-developed empirical literature in economics and finance uses neural networks for forecasting, their potential as a function approximator in rational expectations models has only been recognised more recently.⁴ Maliar et al. (2019) discuss the general applicability of neural networks to computational economics and argue that they can be used to solve many models. Others have shown that equilibrium conditions can be included in the loss function of a neural network to efficiently find global solutions (Azinovic et al., 2022), that neural networks help solve high dimensional models (Fernández-Villaverde et al., 2019), and that they can deal with multicollinear state spaces (Villa and Valaitis, 2022). Our paper demonstrates that neural networks provide an intuitive and plausible resolution to indeterminacy in non-linear models with multiple steady states, adding to well-known results that indeterminacy in linear models leads to instability under learning.

²Iterative e-stability can be thought of as freezing the estimated coefficients for a long period of time, and only re-estimating them at the beginning of the next period.

³Judd et al. (1998).

⁴Cho and Sargent (1996) is an early precursor, using neural nets to represent equilibrium in two repeated games and a dynamic economic model. See Kaastra and Boyd (1996) for a guide to designing neural networks to forecast economic time series.

3 Environment

We keep the analysis as transparent as possible by focussing on an illustrative New Keynesian setup with one state variable (lagged output) and one control variable (inflation). The environment has two outer steady states with a saddlepath structure, plus a central steady state of interest which acts as a “sink” under perfect foresight. In the absence of shocks, the sink steady state is associated with a continuum of convergent paths. We begin by presenting the properties of the system under perfect foresight, as this will prove useful for later comparison with the dynamics of the learnable REE. We then discuss the system’s local determinacy properties under rational expectations.

3.1 New Keynesian model

Inflation π_t and output y_t are determined by a standard New Keynesian Phillips curve, a backward-looking IS curve⁵ and a Taylor Rule with $\phi_\pi < 1$ and $\alpha > 0$:⁶

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa y_t + \epsilon_{\pi,t} \tag{1}$$

$$y_t = \eta y_{t-1} - \sigma(r_t - \mathbb{E}_t \pi_{t+1}) + \epsilon_{y,t} \tag{2}$$

$$r_t = \phi_\pi \pi_t + \alpha \pi_t^3 \tag{3}$$

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 deviation σ_π, σ_y . The Taylor Rule generates multiple convergent paths around the sink steady state, but satisfies the Taylor Principle for more extreme values of inflation as the cubic term bounds the forces that generate multiplicity.

β	κ	η	σ	ϕ_π	α	ρ_π	ρ_y	σ_π	σ_y
0.95	0.05	0.95	0.25	0.5	0.075	0.5	0.5	0.3	0.3

Table 1: Baseline parameter values

The baseline parameterisation in Table 1 is chosen so there are three deterministic steady states in output and inflation, at $(-2, -2), (0, 0)$ and $(2, 2)$. Under perfect foresight, the two outer steady states are local saddles whereas the central steady state is a complex sink with a continuum of convergent paths.

3.2 Perfect foresight solution

The deterministic steady states of the model occur at the intersection of two loci derived from the steady-state versions of the New Keynesian Phillips curve, IS curve and Taylor Rule:

$$\pi = \beta \pi + \kappa y \tag{4}$$

$$y = \eta y - \sigma(\phi_\pi \pi + \alpha \pi^3 - \pi) \tag{5}$$

⁵We choose a backward-looking formulation of the IS curve to have a simple system where agents only need to forecast one variable. This gives a system with one state variable and one control variable that can be visualised in two dimensions.

⁶The key results of our paper are the same with a forward-looking IS curve, even though such a system has no state variable.

Figure 1 shows the locus derived from the Phillips curve as the straight line in black; that derived from the IS curve and Taylor Rule is the s-shape in green. The loci cross at the three deterministic steady states.




Figure 1: New Keynesian model under perfect foresight

To describe the dynamics of the system under perfect foresight, we recognise that $\mathbb{E}_t \pi_{t+1} = \pi_{t+1}$ and abstract from the disturbance terms to express π_t and y_t in terms of π_{t-1} and y_{t-1} :

$$\pi_t = \frac{1}{\beta}(\pi_{t-1} - \kappa y_{t-1}) \quad (6)$$

$$y_t = \frac{\beta}{\beta + \sigma \kappa} \left(\eta y_{t-1} - \sigma \left(\frac{\phi_\pi \beta - 1}{\beta} \pi_t + \alpha \pi_t^3 \right) \right) \quad (7)$$

These dynamics are superimposed on Figure 1 to complete the phase diagram of the system under perfect foresight.⁷ The blue arrows represent paths along which the system will evolve from a given starting point in the y_{t-1}, π_t space of state and control variables. These paths show that the outer steady states are saddles and the steady state of interest is a complex sink.

3.3 Local determinacy properties under rational expectations

It is not possible to characterise the global properties of rational expectations equilibrium in a finite number of parameters, since the model is non-linear and has shocks and state variables (Judd et al., 1998). We can, though, examine the model's local determinacy properties by linearising in the neighbourhood of a steady state. Denoting a steady state of the model by (π^*, y^*) , linearisation around a steady state implies:

$$\hat{\pi}_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} + \kappa \hat{y}_t + \epsilon_{\pi,t} \quad (8)$$

$$\hat{y}_t = \eta \hat{y}_{t-1} - \sigma \left((\phi_\pi + 3\alpha \pi^{*2}) \hat{\pi}_t - \mathbb{E}_t \hat{\pi}_{t+1} \right) + \epsilon_{y,t} \quad (9)$$

⁷The dynamics of the system under perfect foresight are derived from the continuous time analogue of the discrete time system, see Appendix A for details.

The determinacy properties depend on the eigenvalues of this linearised system of equations, following Blanchard and Kahn (1980). At the central steady state $\pi^* = 0$ and the cubic term drops out, so the rational expectations equilibrium is locally indeterminate when $\phi_\pi < 1$. It is straightforward to show that rational expectations equilibrium is determinate in the neighbourhood of the outer steady states.

4 Learning

The learning literature sets rational expectations aside in favour of agents forming expectations without prior knowledge of the nature of the equilibrium. In our case, the key object that agents need to learn is the mapping from lagged variables and the disturbance terms to inflation in the next period, since it is the expectation of next period's inflation that enters the New Keynesian Phillips curve and the IS curve. Evans and Honkapohja (2001) term the mapping a Perceived Law of Motion (PLM), and propose a linear specification which agents update recursively through least squares learning. Learning is said to have converged when agents no longer update their PLM. The non-linearity of our model calls for maximum flexibility in expectations formation, which we respond to by allowing agents to learn using neural networks.⁸ In this section we also discuss our notion of REE, which requires that agents only make unpredictable errors when forming expectations.

4.1 Neural network learning

Neural networks offer a parsimonious functional form which can be estimated efficiently to approximate arbitrarily complex non-linear functions, so are a promising candidate for learning in non-linear models.




Figure 2: Structure of neural network

The neural network our agents use is shown heuristically in Figure 2. It takes a vector of inputs s (lagged variables and the disturbance terms) and transforms it into a vector of outputs \hat{p} (inflation expected next period and current variables) through a hidden layer of nodes. The inputs are passed to the nodes as linear combinations s' of the inputs, with weights w^{hid} and bias terms b^{hid} . A non-linear activation function $\psi(s')$ is then applied to the linear combination entering each node, allowing the input to pass through and

⁸The Parameterised Expectations Algorithm of Den Haan and Marcet (1990) offers a non-linear specification based on power functions that agents update via non-linear least squares learning.

activating the node only if its input is positive, i.e., $\psi(s') = \max(0, s')$. Finally, the output layer aggregates linear combinations of the activated nodes for each output of interest, with weights w^{out} and bias terms b^{out} . This is a single layer feedforward neural network, which can approximate any real-valued continuous function arbitrarily well with any appropriate activation function and a sufficient number of nodes (Cybenko, 1989).

Agents learn the parameters in the neural network using the Levenberg-Marquardt (LM) algorithm.⁹ Starting with initial values for the weights and biases stacked in an $R \times 1$ parameter vector θ_i , the system is simulated for T periods with agents forming their expectations using the neural network. If the neural network has Q outputs of interest then the simulation generates $QT \times 1$ vectors of outcomes p_i , expected outcomes \hat{p}_i and expectational errors $e_i = p_i - \hat{p}_i$. The Jacobian of the expectational errors with respect to the parameter vector is a $QT \times R$ matrix denoted J_i . For the next iteration, the parameter vector is updated to

$$\theta_{i+1} = \theta_i + (J_i' J_i + \mu_i I_R)^{-1} J_i' (p_i - \hat{p}_i) \quad (10)$$

and the system simulated again to obtain $p_{i+1}, \hat{p}_{i+1}, e_{i+1}$ and J_{i+1} . The process continues until the estimate of θ has converged. μ_i is a scalar damping factor that ensures the matrix $(J_i' J_i + \mu_i I_R)$ is positive semi-definite and invertible in early iterations. It is reduced as i increases, meaning that the LM algorithm approaches the Gauss-Newton method as iterations progress. The details are in Hagan and Menhaj (1994), which includes a backpropagation technique that makes calculation of the Jacobian less computationally demanding. We simulate for $T=100,000$ periods between updates, and consider learning as converged when the neural network is able to make very accurate forecasts of current inflation and output.¹⁰

4.2 Verifying 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 on the basis of the information available at the time expectations are formed. If they are we conclude that the equilibrium has the properties of an REE; our notion of REE is hence numerically validated.

The Den Haan and Marcet (1994) 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 for any numerical solution yet the neural network passes it comfortably, giving us confidence that our solution has the properties of an REE.

⁹Alternative training algorithms give virtually identical results, e.g. ADAM (Kingma and Ba, 2015).

¹⁰The accuracy of forecasts is measured by the R^2 between predictions and realisations; a very accurate forecast has an R^2 in excess of 0.99999.

5 Results

Neural network learning converges to an equilibrium that is fundamental, stationary and stable. The central steady state that was locally indeterminate under rational expectations becomes a source under neural network learning, with the system fluctuating in the neighbourhood of one of the locally-determinate steady states until sufficiently large fundamental shocks move it to the neighbourhood of the other locally-determinate steady state. These results are robust to initial conditions and different specifications of the neural network.

5.1 Equilibrium dynamics

The phase diagram of the system under neural network learning is in panel (a) of Figure 3. The two loci in black and green are as before, so the system again has three deterministic steady states. However, the dynamics under neural network learning are very different to those under perfect foresight. The central steady state is now a source that repels the system along a unique path to one of the outer steady states. The path is unlike any of the convergent perfect foresight solutions in Figure 1, nor does it resemble a perturbation around any of the perfect foresight paths.

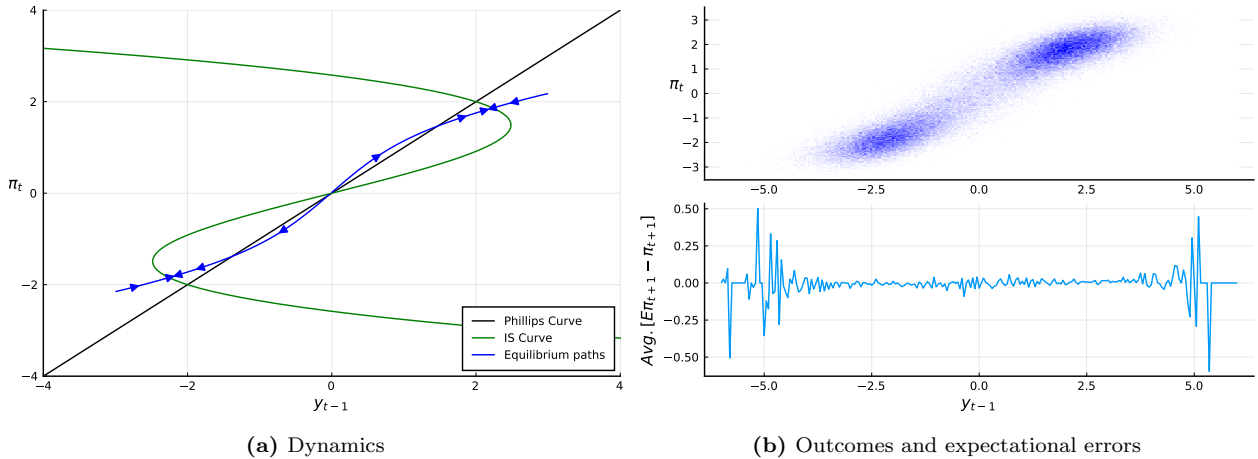


Figure 3: Equilibrium with Neural Network Learning mean dynamics

The upper part of panel (b) in Figure 3 plots the distribution of output and inflation as a heatmap, with darker shades of blue indicating combinations that occur more frequently in simulations of the model. The heatmap shows that simulated outcomes are clustered around the dynamic path in panel (a). As expected given the dynamics, the system spends most of its time in the neighbourhood of one of the outer steady states and it is rare to see observations near the central steady state. Visits to the region near the central steady state are also typically short-lived, since the source dynamics around that steady state rapidly push output and inflation towards one of the stable steady states.

The lower part of panel (b) in Figure 3 shows the average error that agents make when using neural network learning to form their expectation of next period's inflation, as a function of lagged output as the state variable. The errors are not systematically related to the state variable (y_{t-1}), especially when lagged output is close to one of the outer steady states where the system spends most of its time. The unpredictability of the

expectational errors is confirmed by the Den Haan and Marcet (1994) accuracy test results in Table 2, which accept the hypothesis that errors in expectations cannot be predicted even by an extensive set of information available at the time expectations are formed.¹¹ We conclude that dynamics under neural network learning have the characteristics of an REE.

$h(s_t)$	Lower-tail	Upper-tail
Constant	0.0461	0.0520
Constant, y_t and π_t	0.0420	0.0635
Extensive	0.0384	0.0781

Table 2: Accuracy test for neural network learning

Our characterisation of REE has a single equilibrium with multiple steady states. It is fundamental in that its dynamics are entirely driven by the disturbance terms, and stationary because the unconditional distribution of output and inflation does not vary over time. The equilibrium is stable under learning, with small deviations in beliefs disappearing over time as beliefs under neural network learning converge back to values consistent with the REE. A simulation of the model is shown in Figure 4.




Figure 4: Simulation under neural network learning

The behaviour of the system is akin to a regime-switching model, albeit one in which the transitions between “regimes” are governed by the disturbance terms.¹² The system fluctuates around one of the outer steady states until a combination of shocks is sufficient to move it into the basin of attraction of the other steady state. This behaviour suggests that regime-switching specifications could be motivated as reduced forms of nonlinear models, with regions of local indeterminacy acting as sources rather than sinks in REE.

5.2 Learning in small samples

The REE we identify assumes that agents have a sufficiently large number of observations that their beliefs have converged to the limit point of the neural network learning process. To investigate the behaviour of the system in smaller samples, we initialise all outcomes and parameters of the neural network at zero until period

¹¹The extensive information set includes a constant and linear, quadratic and cubic terms in contemporaneous, lagged and twice lagged endogenous variables.

¹²A close analogue is the Filardo (1994) model of Markov-switching with time-varying transition probabilities (TVTP).

100 and then allow agents to start neural network learning using a rolling window of only 400 observations. An example simulation is in Figure 5.




Figure 5: Neural network learning with a small sample

The system approaches the neighbourhood of the upper steady state after about 100 periods of learning, fluctuating around that steady state for roughly 250 periods before moving over to the lower steady state. It subsequently oscillates between periods near each of the outer steady states. We see in Figure 5 that the behaviour of the system in small samples approaches that observed when neural network learning has converged. A large number of observations are needed to pass the Den Haan and Marcet (1994) accuracy test and accept the equilibrium as having the properties of an REE, but only a small number of observations are needed to pin down the equilibrium dynamics and see that the central steady state acts as a source.

5.3 Greater linearity

The cubic term in the Taylor Rule (Eq 3) is a convenient didactic device to generate a region of local indeterminacy in the model. It does though mean that the system is non-linear throughout, even in the neighbourhood of the locally-indeterminate central steady state. We demonstrate that this is not a cause for concern by replicating the analysis and results with a piecewise-linear Taylor Rule (Eq 11). The rule is parameterised with $\phi_\pi = 0.5$ and $\alpha = 0.75$, so the system is locally indeterminate when inflation is in the interval $(-\pi^*, \pi^*)$. When inflation is outside this interval the Taylor Principle is satisfied and the system is locally determinate.

$$r_t = \begin{cases} \phi_\pi \pi_t & -\pi^* \geq \pi_t \geq \pi^* \\ \phi_\pi \pi_t + \alpha(\pi_t - \pi^*) & \pi_t > \pi^* \\ \phi_\pi \pi_t + \alpha(\pi_t + \pi^*) & \pi_t < -\pi^* \end{cases} \quad (11)$$

Dynamics of the model with $\pi^* = 1$ are in Figure 6, where the light blue region indicates the values of inflation for which the system is locally indeterminate. The black line for the steady state of the New Keynesian Phillips curve is as before but the green locus for the combined steady state of the IS curve and the Taylor Rule is now z-shaped. The perfect-foresight dynamics in panel (a) are close to those in Figure 1 for the model with a cubic term in the Taylor Rule. Similarly, the dynamics under neural network learning in panel (b) mirror

those in Figure 3, with the central steady state acting as a source not a sink. We conclude that neural network learning leads to dynamics that are robust to greater linearity around the central steady state.

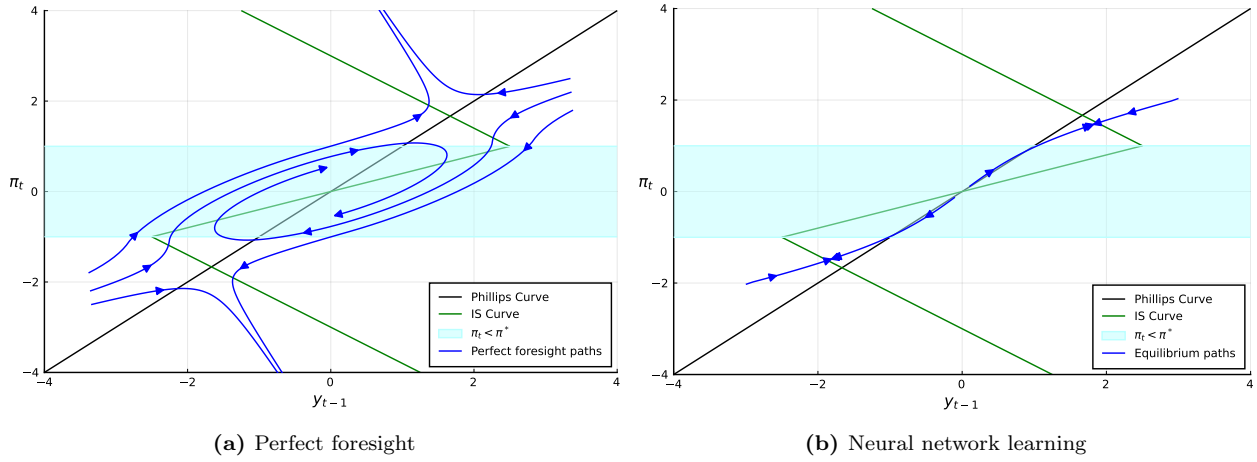


Figure 6: Piecewise linear model with $\pi^* = 1$

The simplicity of the piecewise-linear Taylor Rule prompts us to examine how the size of the region of local indeterminacy affects the nature of the REE identified by neural network learning. Figure 7 for a smaller region shows that neural network learning still uncovers an unstable central steady state and two stable outer steady states, with the unique path under neural network learning in panel (b) closer to linearity. When the region gets larger in Figure 8, neural network learning is unable to distinguish between the separate steady states, instead converging on a unique path that converges to one of the outer steady states.¹³ Whilst in this case learning fails to pick up the central steady state as a source, the system still spends very little time in the neighbourhood of the central steady state.

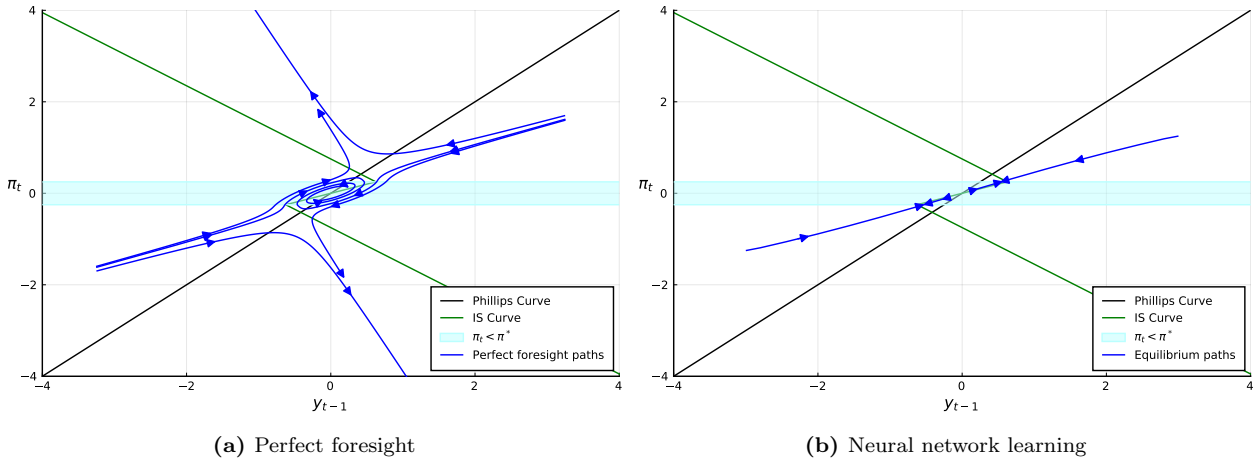


Figure 7: Piecewise linear model with $\pi^* = 0.25$

¹³The unique path converges to the upper steady state in Figure 8, but given the symmetry of the model it is equally likely that neural network learning settles on a path that converges to the lower steady state. A similar convergence to one of the outer steady states also occurs if the variance of the shocks to the fundamental disturbances is very small.

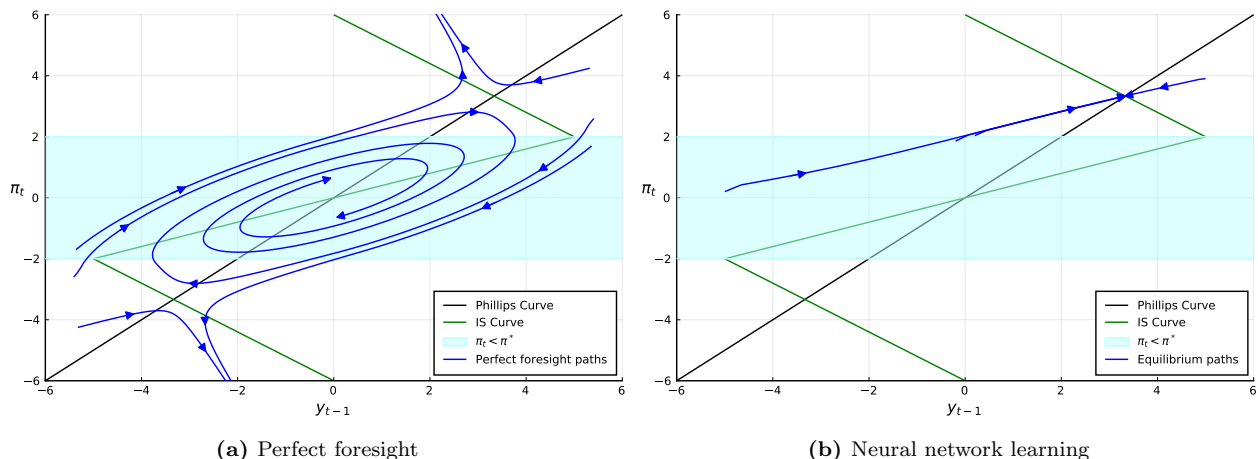


Figure 8: Piecewise linear model with $\pi^* = 2$

5.4 Alternative models


The model analysed so far is tightly parameterised and only supports a limited range of perfect foresight dynamics. A new, more general, model is therefore needed to investigate what happens when agents use neural network learning in a more general setting. In the new model we retain as many features of the original setup as possible, introducing additional degrees of freedom only where necessary to expand the range of possible dynamics. The model continues to have only one state variable y_{t-1} , one control variable p_t , and three deterministic steady states. As before, the outer steady states are locally determinate under rational expectations, but now the properties of the central steady state depend on the parameterisation of the model. In the relevant parameter space the perfect foresight dynamics never have a saddlepath structure, but they can be either a sink or a source with either complex or real eigenvalues.¹⁴

$$p_t = \phi_{p,p} \mathbb{E}_t p_{t+1} + \phi_{p,y} y_t - \alpha y_t^3 + \epsilon_{p,t} \quad (12)$$


$$y_t = \phi_{y,y} y_{t-1} + \phi_{y,p} p_t + \epsilon_{y,t} \quad (13)$$

To cover all possible configurations, it is sufficient to fix $(\phi_{p,p} = 0.95, \phi_{p,y} = 0.5)$ and vary $(\phi_{y,p}, \phi_{y,y})$ such that the central steady state is a complex sink, a real sink, a complex source or a real source. The rows of Figure 9 show an illustrative example of each, with the perfect foresight solution to the left and the REE identified by neural network learning to the right. In the first row the central steady state is a complex sink and dynamics are similar to those in our original model. There are multiple convergent perfect-foresight paths in panel (a) and a unique divergent learnable REE path in panel (b). Subsequent rows show similarly divergent learnable REE paths around the central steady state when it is a real sink (d), complex source (f) or complex sink (h). In none of the cases does the REE identified by neural network learning resemble perfect foresight dynamics.


¹⁴The central steady state can exhibit a saddle path structure for extreme parameterisations of the model, but in that case the outer steady states no longer have saddle path structures. Appendix B presents the conditions for the different cases to occur.




(a) Complex sink: Perfect foresight
 $\phi_{y,p} = 0.1, \phi_{y,y} = 0.9$




(b) Complex sink: REE
 $\phi_{y,p} = 0.1, \phi_{y,y} = 0.9$




(c) Real sink: Perfect foresight
 $\phi_{y,p} = 0.9, \phi_{y,y} = 0.1$




(d) Real sink: REE
 $\phi_{y,p} = 0.9, \phi_{y,y} = 0.1$




(e) Complex source: Perfect foresight
 $\phi_{y,p} = 0.1, \phi_{y,y} = 1.5$



(f) Complex source: REE
 $\phi_{y,p} = 0.1, \phi_{y,y} = 1.5$



(g) Real source: Perfect foresight
 $\phi_{y,p} = 0.1, \phi_{y,y} = 2$



(h) Real source: REE
 $\phi_{y,p} = 0.1, \phi_{y,y} = 2$

Figure 9: Dynamics of alternative models

6 Adaptive learning

This section reduces the sophistication that agents have when forming expectations of future inflation. Instead of giving them access to neural network techniques, we require them to learn using adaptive methods of the type popularised by Evans and Honkapohja (2001)). As we shall show, the dynamics of the resulting equilibrium are qualitatively and quantitatively close to those in the REE identified by neural network learning, most notably in the central steady state always acting as a source from which the system is repelled. This provides further support to the idea that such steady states should be treated as unstable.

6.1 Least squares learning

We assume that agents learn adaptively using the least squares formulation of Evans and Honkapohja (2001), where they update coefficients in Perceived Laws of Motion (PLMs) that describe the evolution of endogenous variables. We consider PLMs that include linear, quadratic and cubic terms in lagged variables:

$$\pi_t = \beta_{\pi,1} + \beta_{\pi,2}y_{t-1} + \beta_{\pi,3}\pi_{t-1} + \beta_{\pi,4}y_{t-1}^2 + \beta_{\pi,5}y_{t-1}^3 + \nu_{\pi,t} \quad (14)$$

$$y_t = \beta_{y,1} + \beta_{y,2}y_{t-1} + \beta_{y,3}\pi_{t-1} + \beta_{y,4}y_{t-1}^2 + \beta_{y,5}y_{t-1}^3 + \nu_{y,t} \quad (15)$$

The agent derives an expectation of inflation in the next period using their PLMs to form expectations $\mathbb{E}_t\pi_t, \mathbb{E}_t\pi_t$ and then rolling the PLM for inflation forward again to obtain $\mathbb{E}_t\pi_{t+1}$.¹⁵ The Actual Laws of Motion (ALMs) for inflation and output then depend on the PLMs and equations (1)-(3).

Agents update their parameter estimates by recursive least squares, with the dependent variable $Y_t = (\pi_t \ y_t)$ and the independent variable $X_t = (1 \ y_{t-1} \ \pi_{t-1} \ y_{t-1}^2 \ y_{t-1}^3)$. The parameter estimates in the 5×2 matrix β_t are updated using the standard formulae:

$$\beta_t = \beta_{t-1} + \frac{1}{t}R_t^{-1}X_t'(Y_t - X_t\beta_t) \quad (16)$$

$$R_t = R_{t-1} + \frac{1}{t}(X_t'X_t - R_t) \quad (17)$$

The convergence and limit properties of recursive least squares learning can be characterised by techniques from stochastic approximation theory. Evans and Honkapohja (2001) show that the first order difference equations (16) and (17) are associated with an ordinary differential equation (ODE) of the form $\frac{d\theta}{d\tau} = h(\theta(\tau))$, where $\theta = \text{vec}(\beta, R)$ is a stacked vector of beliefs and

$$h(\theta(\tau)) = \lim_{t \rightarrow \infty} \mathbb{E} \left[\frac{1}{t} \begin{pmatrix} R_t^{-1}X_t'(Y_t - X_t\beta_t) \\ X_t'X_t - R_t \end{pmatrix} \right] \quad (18)$$

The ODE characterises the mean dynamics of beliefs in notional time τ . Our interest is in whether the ODE has a fixed point $h(\bar{\theta}) = 0$, and if so whether the fixed point is stable for small perturbations. We find that it does, and that the eigenvalues of the Jacobian confirm it as stable under learning. Beliefs therefore converge

¹⁵The results are similar if the agent has a single PLM that specifies inflation as a function of the second lag of inflation and output, in which case inflation expectations can be read off the PLM directly. They are also much the same if agents observe the disturbance terms and incorporate them in their PLMs.

in the associated ODE, and parameter estimates converge under least squares learning.

6.2 Equilibrium dynamics

Recursive least squares learning converges, as did neural network learning, to an equilibrium that is fundamental, stationary and stable. The fixed point of the mean dynamics picks up the negative system feedback via the cubic terms in the PLMs and puts no weight on constant and quadratic terms:

$$\bar{\beta} = \begin{pmatrix} 0.00 & 1.85 & 0.09 & 0.00 & -0.21 \\ 0.00 & 1.07 & 0.01 & 0.00 & -0.02 \end{pmatrix} \quad (19)$$

The equilibrium under adaptive learning in panel (a) of Figure 10 has a phase diagram that is qualitatively and quantitatively close to that in Figure 3 for the REE identified by neural network learning. The central steady state is again a source and the outer steady states are stable.

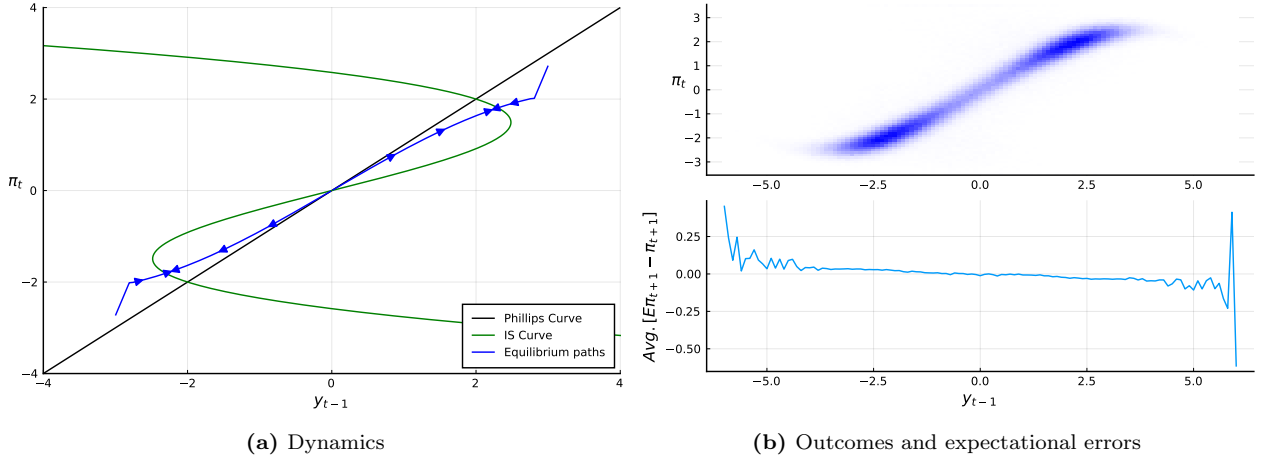


Figure 10: New Keynesian model under recursive least squares learning

The heatmap in the upper part of panel (b) in Figure 1 confirms that the simulated system spends most of its time near to one of the outer steady states, with only limited forays into the neighbourhood of the central steady state. The lower part of panel (b) in Figure 1 demonstrates that agents make small yet systematic errors in expectations even when adaptive learning has converged, which is not surprising given they are less sophisticated than neural network learners. The errors in expectations fail the Den Haan and Marcet (1994) accuracy test by a large margin, so are predictable and the equilibrium under adaptive learning lacks the characteristics of an REE. Its key features do though mimic those of the REE, giving us further confidence that we have identified the most plausible equilibrium dynamics.

7 Empirical tests for indeterminacy

Lubik and Schorfheide (2004) propose a test for indeterminacy and sunspots in likelihood-based estimation of dynamic general equilibrium models. It compares the probability that data is generated by a locally determinate system as opposed to being perturbations of perfect foresight paths in a locally indeterminate system. In

this section, we apply the test to data simulated by our model. Since REE dynamics in the neighbourhood of our model's central steady state resemble neither a saddlepath stable system nor perturbations to perfect foresight paths, we investigate whether the test may misleadingly favour sunspots as a driver of the data.

7.1 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 .¹⁶ 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 Eq 20 and applying the Jordan decomposition $A = P^{-1} \Lambda P$:

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

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

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

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

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 Eqs 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 Eqs 22 and 23 do not have a unique solution. In this case, Farmer et al. (2015) suggests 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:¹⁷

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

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

¹⁶Blanchard and Kahn (1980) or Sims (2002).

¹⁷Lubik and Schorfheide (2004) explicitly model a dependence of the sunspot on fundamental shocks.

7.2 Estimation results

The Lubik and Schorfheide (2004) test involves estimating determinate and indeterminate models on the same data, and comparing their relative fit through the posterior odds ratio. We perform the test on 10,000 periods of simulated data from our learnable REE.

Parameter	Prior mean	Posterior mean	
		Determinate	Indeterminate
β	$\mathcal{N}(1, 1)$	-0.34 (-0.61, 0.09)	1.06 (0.97, 1.12)
κ	$\mathcal{N}(0, 1)$	1.02 (0.83, 1.23)	-0.03 (-0.08, 0.04)
η	$\mathcal{N}(1, 1)$	0.95 (0.89, 1.03)	0.85 (0.82, 0.88)
σ	$\mathcal{N}(0, 1)$	-0.03 (-0.36, 0.20)	0.41 (0.32, 0.52)
ϕ_π	$\mathcal{N}(2, 1)$	0.57 (-0.32, 1.18)	0.55 (0.42, 0.70)
ρ_π	$\mathcal{N}(0, 1)$	0.59 (0.58, 0.60)	0.45 (0.41, 0.49)
ρ_y	$\mathcal{N}(0, 1)$	0.48 (0.46, 0.50)	0.50 (0.48, 0.51)
σ_π	$\mathcal{IG}(0.5, 2)$	0.48 (0.42, 0.55)	0.25 (0.21, 0.28)
σ_y	$\mathcal{IG}(0.5, 2)$	0.21 (0.18, 0.26)	0.17 (0.17, 0.18)
σ_ζ	$\mathcal{IG}(0.5, 2)$	-	0.39 (0.37, 0.42)
$\omega_{\pi, \zeta}$	$\mathcal{B}(0, 0.3, -1, 1)$	-	0.70 (0.61, 0.82)
$\omega_{y, \zeta}$	$\mathcal{B}(0, 0.3, -1, 1)$	-	0.51 (0.43, 0.60)
Log data density		-3927	-3889

Table 3: Bayesian estimation results

The Bayesian estimation results are in Table 3.¹⁸ They show a clear preference for the sunspot model. The log posterior for the determinate model is -3927 whereas for indeterminate model it is -3889, so the posterior odds ratio overwhelmingly favours indeterminacy and the simulated data strongly supports 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 significantly more volatile than either of them.

The positive correlations between the sunspot shock and innovations to the disturbance terms are consistent with the central steady state in our REE being a local source. Consider the model at the central steady state and small positive innovations $\epsilon_{\pi, t}, \epsilon_{y, t} > 0$ that send inflation and output to $\pi_t, y_t > 0$. At this point, the source nature of REE dynamics propels inflation further away from the steady state and $E_t \pi_{t+1} > \pi_t, E_t y_{t+1} > y_t$. This divergent behaviour is difficult for the linearised indeterminate model to explain, since the equilibrium under rational expectations is a real sink.¹⁹ However, there are perfect-foresight paths that temporarily imply rising inflation or output. The positive correlation of the sunspot shock picks out such paths, and helps the indeterminate model explain why inflation and output continue to rise after the initial innovations.

¹⁸The models are estimated in Dynare using Metropolis Hastings MCMC, as described in Guerrón-Quintana and Nason (2013).

¹⁹It is even more difficult for the linearised determinate system to explain.

8 Application

This section applies neural network learning to study the characteristics of the learnable REE in a New Keynesian model with an effective lower bound on the nominal interest rate. Such an environment has been of renewed interest since Benhabib et al. (2001), who pointed out the existence of a second deflationary steady state in such a model. The second steady state is indeterminate but not stable under learning,²⁰ prompting Evans et al. (2022) and Eggertsson et al. (2019) to argue for the introduction of additional bounds on the region of indeterminacy. The resulting systems have three steady states, the outer ones determinate and the central one indeterminate, so map directly into our setup.

8.1 A bounded New Keynesian model

We assume that inflation, output and the nominal interest rate are all subject to lower bounds, with those on inflation and output capturing the downward price-setting rigidity and non-negativity of consumption in the microfounded model of Evans et al. (2022). The system is otherwise as before, except for the IS curve (Eq 26) being both backward and forward looking and the Taylor Rule (Eq 27) being defined over inflation and output. The New Keynesian Phillips Curve (Eq 25) is unchanged.

$$\pi_t = \max\{\beta\mathbb{E}_t\pi_{t+1} + \kappa y_t, \pi^{lim}\} + \epsilon_{\pi,t} \quad (25)$$

$$y_t = \max\{(1 - \eta)\mathbb{E}_t y_{t+1} + \eta y_{t-1} - \sigma(r_t - \mathbb{E}_t \pi_{t+1}), y^{lim}\} + \epsilon_{y,t} \quad (26)$$

$$r_t = \max\{\phi_\pi \pi_t + \phi_y y_t, r^{lim}\} \quad (27)$$

The disturbance terms in Eq 25 and Eq 26 follow independent AR(1) processes, with innovations in normal times that have a standard deviation of $\sqrt{0.001}$. On rare occasions the innovations are much larger, following the “disaster shocks” of Gourio (2012). These large innovations occur on average every 100 periods, and have a magnitude that in the metric of normal times would correspond to ± 10 standard deviation shocks. The model’s other parameters are calibrated in Table 4. The calibration of the lower bounds implies a lower steady state in which the nominal interest rate is 2% below its upper steady state value, with inflation and output 3% and 5% below their upper steady state values.

β	κ	η	σ	ϕ_π	ϕ_y	ρ_π	ρ_y	r^{lim}	y^{lim}	π^{lim}
0.95	0.1	0.5	0.25	1.5	0.5	0.5	0.5	-0.02	-0.05	-0.03

Table 4: Parameter values

8.2 Results

The steady state loci of the model are shown in both panels of Figure 11, where the blue shaded region shows combinations of output and inflation for which the nominal interest rate is constrained by the effective lower bound. The black locus for the New Keynesian Phillips Curve has a kink at the lower bound on inflation and the green locus that combines the IS curve and Taylor Rule has two kinks, one at the lower bound for output and one where inflation is so low that the lower bound on the nominal interest rate constrains the Taylor Rule.

²⁰Christiano et al. (2018).

The loci cross three times, giving lower, central and upper steady states as in Figures 1 and 6 for the original and piecewise linear models.

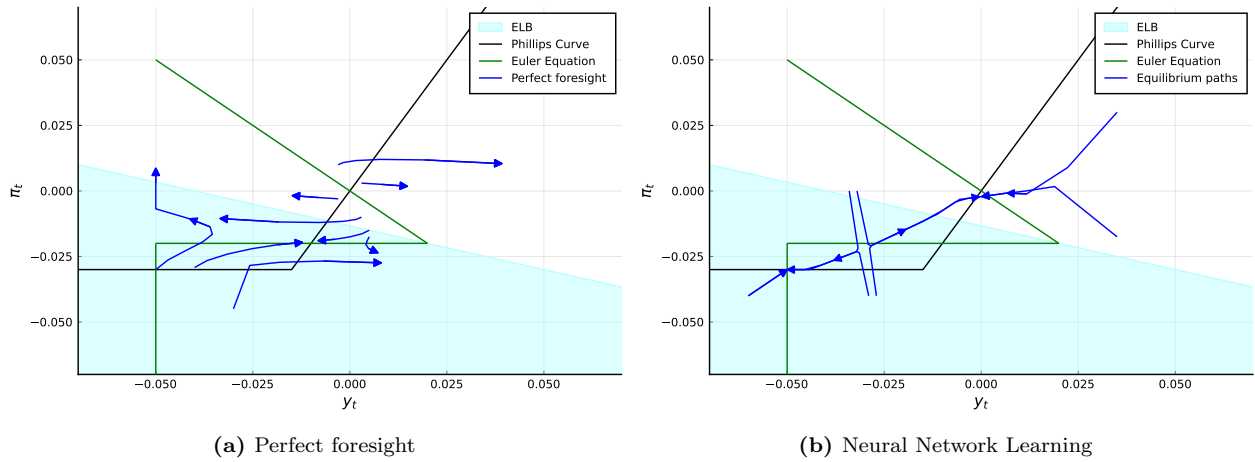


Figure 11: Bounded New Keynesian model with Zero Lower Bound

The perfect foresight dynamics in panel (a) of Figure 11 reflect the local indeterminacy of the central steady state and the local determinacy of the outer steady states. The REE dynamics under neural network learning in panel (b) show the now familiar pattern, where the central steady state acts as a source from which there is a unique divergent path to the outer stable steady states. The system spends most of its time in simulations close to the outer steady states, moving between them only when there is a suitable “disaster shock”. It is rare to have observations in the neighbourhood of the locally-indeterminate central steady state.

9 Closing remarks

Multiple equilibria are common in environments where agents need to make forecasts. In particular, many papers have identified economic forces that give rise to a locally-indeterminate steady state surrounded by a continuum of convergent perfect foresight paths. We argue that such steady states should be treated as unstable under rational expectations, a conclusion we reached by looking for REE that are learnable. Even when agents are endowed with sophisticated learning capacities, we show that agents never learn equilibria that resemble perturbations of convergent perfect foresight paths around a locally-indeterminate steady state. Rather, the learnable REE is associated with a dynamic path which repels away from the steady state. We also showed that an econometrician confronted with data reflecting this instability may nevertheless improperly infer that behaviour around this steady state is convergent and affected by sunspots.

Our call to treat a locally-indeterminate steady state as unstable has parallels in the work on coordination failures of Cooper and John (1988). They show that strategic complementarities are necessary to generate multiple equilibria, and that static models with localised strategic complementarity often give rise to three steady states. The two outer steady states are typically stable under tâtonnement and the central steady state is unstable, just as in the learnable REE of our dynamic environment. In neither case would we expect the system to spend significant amounts of time in the neighbourhood of the unstable central steady state.

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.
- Jasmina Arifovic, James Bullard, and Olena Kostyshnya. Social learning and monetary policy rules. *Economic Journal*, 123(567):38–76, 2013.
- Marlon Azinovic, Luca Gaegauf, and Simon Scheidegger. Deep Equilibrium Nets. *International Economic Review*, 63(4):1471–1525, 2022.
- Jess Benhabib and Roger EA Farmer. Indeterminacy and Increasing Returns. *Journal of Economic Theory*, 63(1):19–41, 1994.
- Jess Benhabib, Stephanie Schmitt-Grohé, and Martin Uribe. The Perils of Taylor Rules. *Journal of Economic Theory*, 96(1-2):40–69, 2001.
- Michele Berardi and John Duffy. Real-Time, Adaptive Learning via Parameterized Expectations. *Macroeconomic Dynamics*, 19(2):245–269, 2015.
- 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.
- James Bullard. Learning Equilibria. *Journal of Economic Theory*, 64(2):468–485, 1994.
- 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.
- In-Koo Cho and Thomas J Sargent. Neural networks for encoding and adapting in dynamic economies. In *Handbook of Computational Economics*, volume 1, pages 441–470. Elsevier, 1996.
- 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.
- Richard Clarida, Jordi Gali, and Mark Gertler. Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory. *Quarterly Journal of Economics*, 115(1):147–180, 2000.
- Russell Cooper and Andrew John. Coordinating Coordination Failures in Keynesian Models. *Quarterly Journal of Economics*, 103(3):441–463, 1988.
- 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. Solving the Stochastic Growth Model by Parameterizing Expectations. *Journal of Business & Economic Statistics*, 8(1):31–34, 1990.
- 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, 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 Evans and Bruce McGough. Monetary policy, indeterminacy and learning. *Journal of Economic Dynamics and Control*, 29(11):1809–1840, 2005.
- George W Evans and Seppo Honkapohja. *Learning and Expectations in Macroeconomics*. Princeton University Press, 2001.
- 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.
- Jesús Fernández-Villaverde, Samuel Hurtado, and Galo Nuno. Financial Frictions and the Wealth Distribution. NBER Working Papers 26302, National Bureau of Economic Research, 2019.
- Jesus Fernandez-Villaverde, Federico Mandelman, Yang Yu, and Francesco Zanetti. Search Complementarities, Aggregate Fluctuations, and Fiscal Policy. NBER Working Papers 26210, National Bureau of Economic Research, 2020.
- Andrew J Filardo. Business-Cycle Phases and Their Transitional Dynamics. *Journal of Business & Economic Statistics*, 12(3):299–308, 1994.
- Jordi Gali. Product diversity, endogenous markups, and development traps. *Journal of Monetary Economics*, 36(1):39–63, 1995.
- Francois Gourio. Disaster Risk and Business Cycles. *American Economic Review*, 102(6):2734–66, 2012.
- 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.
- 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.
- Martin T Hagan and Mohammad B Menhaj. Training feedforward networks with the Marquardt algorithm. *IEEE transactions on Neural Networks*, 5(6):989–993, 1994.
- Cars Hommes and Gerhard Sorger. Consistent Expectations Equilibria. *Macroeconomic Dynamics*, 2(3):287–321, 1998.

- Zhen Huo and José-Víctor Ríos-Rull. Paradox of Thrift Recessions. NBER Working Papers 19443, National Bureau of Economic Research, 2013.
- Kenneth L Judd et al. Numerical Methods in Economics. *MIT Press Books*, 1, 1998.
- Ieabeling Kaastra and Milton Boyd. Designing a neural network for forecasting financial and economic time series. *Neurocomputing*, 10(3):215–236, 1996.
- Greg Kaplan and Guido Menzio. Shopping Externalities and Self-Fulfilling Unemployment Fluctuations. *Journal of Political Economy*, 124(3):771–825, 2016.
- Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- Paul Krugman. History Versus Expectations. *Quarterly Journal of Economics*, 106(2):651–667, 1991.
- Thomas A Lubik and Frank Schorfheide. Testing for Indeterminacy: An Application to U.S. Monetary Policy. *American Economic Review*, 94(1):190–217, 2004.
- Robert Lucas. Adaptive Behavior and Economic Theory. *Journal of Business*, 59(4):S401–26, 1986.
- 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. Will Artificial Intelligence Replace Computational Economists Any Time Soon? CEPR Discussion Papers 14024, Centre for Economic Policy Research, 2019.
- 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.
- Bennett T McCallum. E-stability vis-a-vis determinacy results for a broad class of linear rational expectations models. *Journal of Economic Dynamics and Control*, 31(4):1376–1391, 2007.
- Christopher A Sims. Solving Linear Rational Expectations Models. *Computational Economics*, 20(1):1–20, 2002.
- Alessandro T Villa and Vytautas Valaitis. A Machine Learning Projection Methods for Macro-Finance Models. FRB of Chicago Working Papers 2022-19, Federal Reserve Bank of Chicago, 2022.
- Michael Woodford. Learning to Believe in Sunspots. *Econometrica*, 58(2):277–307, 1990.
- Michael Woodford. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, 2003.

A Continuous time analogue of the discrete time model

The discrete time perfect foresight dynamics (Eq 6) and (Eq 7) can be interpreted as a discretisation of an underlying continuous time process, where the time step of the discretisation is $T = 1$. To obtain the phase diagram in Figure 1, we discretise the continuous time dynamics at smaller time steps. The general formula for discretisation with any time step T implies:

$$\pi_{t+k} = e^{AT} \pi_t + A^{-1}(e^{AT} - 1)B y_t \quad (28)$$

$$y_{t+k} = e^{CT} y_t + C^{-1}(e^{CT} - 1)D \left(\frac{\phi_\pi \beta - 1}{\beta} \pi_t + \alpha \pi_t^3 \right) \quad (29)$$

$$e^A = \frac{1}{\beta} \quad (30)$$

$$A^{-1}(e^A - 1)B = -\frac{\kappa}{\beta} \quad (31)$$

$$e^C = \frac{\beta \eta}{\beta + \sigma \kappa} \quad (32)$$

$$C^{-1}(e^C - 1)D = \frac{\beta \sigma}{\beta + \sigma \kappa} \quad (33)$$

B Dynamics in alternative models

The perfect foresight dynamics are obtained from Eqs 12 and 13, with $\epsilon_{p,t} = 0, \epsilon_{y,t} = 0, \mathbb{E}_t p_{t+1} = p_{t+1}$.

$$p_t = \phi_{p,p} p_{t+1} + \phi_{p,y} y_t - \alpha y_t^3 \quad (34)$$

$$y_t = \phi_{y,y} y_{t-1} + \phi_{y,p} p_t \quad (35)$$

The deterministic steady states satisfy the two black loci in each panel of Figure 9.

$$(1 - \phi_{p,p}) p^* = \phi_{p,y} y^* - \alpha y^{*3} \quad (36)$$

$$(1 - \phi_{y,y}) y^* = \phi_{y,p} \phi_{p,p} p^* \quad (37)$$

Linearising around a deterministic steady state implies

$$\hat{p}_t = \phi_{p,p} \mathbb{E}_t \hat{p}_{t+1} + (\phi_{p,y} - 3y^{*2} \alpha) \hat{y}_t + \epsilon_{p,t} \quad (38)$$

$$\hat{y}_t = \phi_{y,y} \hat{y}_{t-1} + \phi_{y,p} \hat{p}_t + \epsilon_{y,t} \quad (39)$$

and defining $\hat{x}_t = (\hat{p}_t \hat{y}_{t-1})'$ and $\epsilon_t = (\epsilon_{p,t} \epsilon_{y,t})$ admits a state space representation of the model:

$$\mathbb{E}_t \hat{x}_{t+1} = \Phi \hat{x}_t + \Psi \epsilon_t \quad (40)$$

$$\Phi = \begin{pmatrix} \frac{1 - (\phi_{p,y} - 3y^* \alpha) \phi_{y,p}}{\phi_{p,p}} & -\frac{(\phi_{p,y} - 3y^* \alpha) \phi_{y,y}}{\phi_{p,p}} \\ \phi_{y,p} & \phi_{y,y} \end{pmatrix} \quad (41)$$

$$\Psi = \begin{pmatrix} -\frac{1}{\phi_{p,p}} & -\frac{\phi_{p,y} - 3y^* \alpha}{\phi_{p,p}} \\ 0 & 1 \end{pmatrix} \quad (42)$$

The dynamics in the neighbourhood of a deterministic steady state are determined by the eigenvalues of the coefficient matrix Φ , which can be expressed in terms of the parameters $\phi_{p,p}$, $\phi_{p,y}$, $\phi_{y,p}$ and $\phi_{y,y}$ using the formula $\lambda = tr(\Phi) \pm \sqrt{tr(\Phi)^2/4 - \det(\Phi)}$. The cubic term with $\alpha > 0$ in Eq 12 ensures that the two outer steady states are always determinate.²¹ Figure 12 shows the characteristics of the central steady state as a function of $\phi_{y,y}$ and $\phi_{y,p}$, for fixed $\phi_{p,p} = 0.95$ and $\phi_{p,y} = 0.5$.




Figure 12: Properties of the central steady state when $\phi_{p,p} = 0.95$, $\phi_{p,y} = 0.5$

²¹Proof on request.