Social learning and expectational stability*

George Evans
University of Oregon
University of St Andrews

Bruce McGough
University of Oregon

September 22nd 2023

Abstract

Stability features of social learning (SL) dynamics are examined. We show SL can be formulated as a stochastic recursive algorithm, making it possible to analyze asymptotics using the familiar differential-equation approach. For a simple univariate model, this approach reduces to the E-stability principle, though in prominent instability cases divergence is *exceedingly* slow compared to adaptive learning (AL). We locate differing fitness criteria as the source of the slower evolution rates of SL compared to AL. Modified AL and SL learning dynamics models are developed and used to illustrate the different implications of policy change in a standard New Keynesian model. We anticipate that the central question going forward will be how best to combine the two approaches when modeling adaptation to structural change.

JEL Classifications: E31; E32; E52; D84; D83
Key Words: Social learning, E-stability, bounded rationality, heterogeneous expectations

*An earlier version of this paper was presented at *February 23, 2023, the Simon Fraser University Computational and Experimental Workshop in honour of Jasmina Arifovic*. We thank Jungang Li for excellent research assistance and participants at University of Oregon and the University of St. Andrews Macro Groups for comments.

1 Introduction

The rational expectations (RE) approach to expectation formation is subject to two lines of critique: it typically presumes an unrealistic degree of information about the economy; and, since in macroeconomic contexts it is comes as part of a *rational expectations equilibrium* (REE), it begs a mechanism by which the equilibrium will be attained. This last point is particularly salient in models with multiple REE.

One natural reaction to this critique is to replace RE with plausible behavioral rules, which may, of course, lead to deviations from RE. A prominent example is adaptive learning (AL), which can be viewed as a bounded rationality approach that attempts to keep close to the spirit of RE in that agents should not repeatedly make systematic forecasting errors. Learning agents are modeled as following an adaptive process for revising forecast rules that can over time eliminate predictable mistakes, and possibly lead to convergence, appropriately defined, to RE. Sargent (2008) calls this an adaptive evolutionary approach, and a widely used implementation is to model agents as statisticians or econometricians who update their forecast rules over time in accordance with (recursive) least-squares learning.[1] This "adaptive learning" (AL) approach has been extensively used in macroeconomics in theoretical, empirical and in policy-oriented settings. For a recent survey see Evans and McGough (2020).

There are other adaptive evolutionary models, with the potential to converge to RE, that rely on non-econometric boundedly rational approaches.[2] A particularly prominent example is the "genetic-algorithm learning" or "social learning" (SL) approach pioneered by Jasmina Arifovic. SL emphasizes heterogeneity of expectations across agents. Early papers include applications to cobweb, hyperinflation, exchange-rate and growth transition models.[3] Recent applications to questions concerning monetary policy include Arifovic, Bullard, and Kostyshyna (2013), Arifovic, Schmitt-Grohe, and Uribe (2018), and Arifovic, Grimaud, Salle, and Vermandel (2020). Major attractions of SL are that it embraces the heterogeneity of expectations that is so evident in survey data, and that the mechanisms by which forecasts adapt can be interpreted as experimentation, interchange between agents, and an evolutionary direction toward more successful forecasting rules. The potential importance of social learning following large unanticipated shocks has recently been emphasized by Bullard (2023), who argues that social learning dynamics, which can amplify the inertia of beliefs, provide an explanation for recent rapid recoveries following severe downturns, e.g. the Great Recession and the Covid Pandemic.

While SL has these major benefits, a disadvantage to date of SL is that there are few, if any, theoretical results on its asymptotic properties. By contrast, there is a large and well-established literature on the asymptotic stability properties of REE, and more generally *restricted perceptions*

[1] See, for example, Bray and Savin (1986), Marcet and Sargent (1989), Evans (1989), Woodford (1990) and Evans and Honkapohja (2001). See also Evans (1985).

[2] For an early survey of alternative approaches see Sargent (1993).

[3] Arifovic (1994), Arifovic (1995), Arifovic (1996) and Arifovic, Bullard, and Duffy (1997).

equilibria (RPE), under adaptive learning.[4] A particularly prominent estimation procedure used in this literature is *recursive least-squares learning:* see e.g. Bray and Savin (1986), Marcet and Sargent (1989) and Evans and Honkapohja (2001). Convergence results can be obtained using the theory of stochastic recursive algorithms (SRA), introduced by Robbins and Monro (1951), and generalized by e.g. Ljung (1977), Kushner and Clark (1978) and Benveniste, Metivier, and Priouret (1990).

Asymptotic analysis of SRAs is commonly conducted using the *ODE method*: See Benveniste, Metivier, and Priouret (1990), Chapter 2, for a general introduction. This method assigns to the SRA an ordinary differential equation – sometimes referred to in the AL literature as the *mean dynamics* – that locally approximates the expected trajectory of the algorithm. Using these mean dynamics, Evans and Honkapohja (2001) developed the *E-stability principle*, which states that the stability of a central and easily computed component of the mean dynamics governs local stability of an REE/RPE under AL and closely related learning rules. This principle has been validated even in cases not covered by standard SRA assumptions. Clearly it would be desirable to develop a theoretical connection between SL and E-stabilty.

In addition to lacking a theoretical basis for stability analysis, computational results in particular models suggest discrepancies between the stability properties of SL and AL. This is of considerable importance theoretically, but also, as illustrated in Section 4.2, there can be significant policy implications. A prominent example of the importance of stability under learning concerns contemporaneous "Taylor-type" interest-rate rules used by Central Banks. For the benchmark New Keynesian model, Bullard and Mitra (2002) showed that E-stability of the REE requires active monetary policy: passive policies lead to instability. However, in the same model, and for the same calibration, Arifovic, Bullard, and Kostyshyna (2013) provided striking computational evidence that under SL the RE steady state is stable *also* under passive policy. This provides an example of an REE that is not E-stable but that appears clearly to be stable under SL. This example calls into question the generality of the E-stability principle, and it raises a further question: what features of SL generate apparent stability of an REE that is unstable under AL?

The current paper makes progress on these issues. To do so we construct a univariate theoretical laboratory in which we can examine SL and compare it to AL. A major contribution of our paper is that we represent SL as an SRA amenable to the ODE method. This can be viewed as providing a proof of concept that social learning can be analyzed with the same tools as used to study adaptive learning. Within our framework, and with a few simplifying assumptions, we derive the model's mean dynamics. We find the asymptotic stability properties of SL and AL are *the same*, and both are governed by the E-stability principle. How can these theoretical results be reconciled with the simulation results?

We show that under SL the instability can be *exceedingly* slow to emerge, a phenomenon we call "stable instability." The other main contribution of this paper is to track down the source,

[4]RPE is a generalization of REE in which forecast rules are optimized within a restricted class under consideration. See Evans and Honkapohja (2001), Branch (2006), and Evans and McGough (2020).

in these cases, of the very different simulation results of SL and AL, over realistic horizons, in the case of stable instability. We show that the reasons for the different results lie in the fitness measure used under SL in the "tournament" stage, and the "gain sequences" used under AL to update forecasts. To demonstrate this we develop a modified AL that mimics SL tournament play, and a modified SL using a fitness measure that mimics standard AL updating.

We examine the policy implications of these results within the standard bivariate New Keynesian model studied by Arifovic, Bullard, and Kostyshyna (2013). In line with the simulation results of Arifovic, Bullard, and Kostyshyna (2013) the (minimum state variable) REE exhibits stable instability: instability under SL cannot be detected even in extremely long simulations. However, SL also has surprising implications for policy change. Specifically, suppose the central bank (CB) has an inflation target $\pi^* = 5\%$, that the CB follows a Taylor-type interest rate rule based on that target, and that the economy is initially in a steady state with $\pi = \pi^* = 5\%$. Suppose policymakers then decide to reduce their target inflation rate to $\pi^* = 0$, which of course implies a corresponding reduction in the steady state nominal interest rate i. Assuming an active interest-rate rule, under AL the implications are conventional: the CB increases interest rates because inflation is above the new zero inflation target. This leads to higher ex-ante real interest rates r, which reduces output and inflation en route to the new steady state.

Under SL the implications of the reduction in the π^* target are unusual and depend on whether active or passive policy is followed. Under an active policy rule with SL there is a persistent drop in output and the reduction in inflation is imperceptibly small. If instead a passive policy rule is followed with SL, there is a persistent increase in output, while again changes in the inflation rate are imperceptibly small. Finally, if we consider the modified SL model, using a fitness measure more aligned to AL, the results under active policy are quite similar to the AL results.

The paper is organized as follows. Section 2 reviews and distinguishes the standard adaptive learning and social learning frameworks, and covers the E-stability principle. Section 3 analyzes AL and SL in a univariate laboratory. The E-stability is shown to govern asymptotic dynamics of both learning mechanisms, and additional results on constant gain dynamics are provided. The concept of stable instability is introduced, and examined via numerical analysis of modified learning algorithms. Section 4 considers the policy implications of stable instability in a New Keynesian framework. Section 5 concludes.

2 Review of adaptive learning and social learning

We consider learning dynamics in linear macro models featuring expectational feedback in the form of a dependence on average forecasts across a finite number of agents. For the purposes of discussion, we take our model to have the form

$$y_t = \beta M^{-1} \sum_{i=1}^{M} E_t(i) y_{t+1} + v_t, \tag{1}$$

though nothing in principle precludes the inclusion of lags. Here M is the number of agents (taken to be even for convenience), $y_t \in \mathbb{R}^n$ is the endogenous state, β is an $n \times n$ matrix, and $v_t \in \mathbb{R}^n$ is an exogenous, zero-mean, iid process, though our results can be easily extended to allow for serial correlation, provided stationarity is imposed. Finally, $E_t(i)\, y_{t+1}$ is the forecast of y_{t+1} made by agent i in period t. If agents hold rational expectations then there is a unique *minimal state variable* (MSV) solution, which takes the form $y_t = v_t$.

2.1 Adaptive learning

In the MSV solution, under RE $E_t(i)\, y_{t+1} = 0$ for all agents i and all times t, though the rational expectations hypothesis is silent on how agents came to coordinate on this forecast. By contrast, the adaptive learning literature adopts the view that agents estimate their forecast models, updating them over time as new data become available. In this way, the model's MSV solution may (or may not) identify a plausible emergent outcome for the economy.

To provide an explicit implementation of adaptive learning, we assume that each boundedly rational agent uses a forecasting model that is consistent with the MSV solution, i.e. they regress on a constant. Notationally, for $1 \le k \le n$ we let $\phi_{kt}(i)$ represent the forecast model used by agent i in period t to forecast y_{kt+1}, so that $E_t(i)\, y_{kt+1} = \phi_{kt}(i)$. The column vector $\phi_t(i) = (\phi_{1t}(i), \dots, \phi_{nt}(i))^{\mathsf{T}}$ will sometimes be referred to as the period t *perceived law of motion* (PLM) of agent i. Note in particular that $\phi_t(i) \in \mathbb{R}^n$ identifies agent i's time t beliefs.

Agent i uses *recursive least squares* (RLS) to update beliefs, which amounts to computing the (possibly weighted) sample mean:

$$\phi_{t+1}(i) = \phi_t(i) + \gamma_{t+1}\,(y_t - \phi_t(i)). \tag{2}$$

Here γ_t is the gain sequence, which measures the weight placed on new information.[5] In general, γ_t can be taken as decreasing to zero or as a positive constant. In this latter case, the algorithm is referred to as *constant gain learning* (CGL). With $\gamma_t = t^{-1}$, $\phi_t(i)$ is the usual sample mean.

The agents' forecast models $\phi_t(i)$ determine the *actual law of motion* (ALM) for y_t:

$$y_t = \beta M^{-1} \sum_{i=1}^{M} \phi_t(i) + v_t. \tag{3}$$

Equations (2) and (3) can be combined to determine a stochastic recursive system in the vector $\phi_t = (\phi_t(i))_{i=1}^{M} \in \mathbb{R}^{n \cdot M}$. The stability question is: under what conditions do the $\phi_t(i)$ converge (in an appropriate sense) to zero?

[5] Agent-specific gains have also been examined (see Evans, Honkapohja, and Marimon (2001)), and would not play a prominent role in our analysis, so for simplicity we assume all agents the same gain.

2.2 Social learning

Under adaptive learning, agents modify their beliefs by incrementally accounting for the most recent forecast error. Social learning also uses forecast performance as the measure of success, but agents' beliefs are modified over time in a manner analogous to the evolution of genes – via mixing, mutation, and natural selection.

To understand social learning dynamics, it is helpful to lean heavily on the biology metaphor: the period t PLMs $\phi_{kt}(i)$ are *genes*, with the PLMs of a given agent, $\phi_t(i)$, identifying a period t *chromosome*. Somewhat awkwardly, the population of period t chromosomes ϕ_t comprises the *gene pool.*

The gene pool evolves over time via *cross-over*, *mutation*, and *tournament selection* (in that order), as described in detail below. It is useful to view each of these mechanisms as an operator acting on the gene pool.

1. *Crossover.* Informally, agents are randomly paired, and pairs swap some of their genes. Formally, let x be uniformly distributed over $\{0,1\}$. Each period $M/2$ pairs of chromosomes are randomly selected *without replacement* (this is why we assume there is an even number of agents). With probability p_c, pairs engage in *crossover,* which means they swap gene k with probability $1/2$. That is, a sample $\{x_1,\ldots,x_n\}$ of independent draws of x is obtained, and the matched agents swap their k^{th}-PLMs provided $x_k = 1$. Thus if, in period t, agents i and j engage in crossover, and if $x_k = 1$, then $\hat{\phi}_{kt}(j) = \phi_{kt}(i)$ and $\hat{\phi}_{kt}(i) = \phi_{kt}(j)$, where the hat identifies the modified chromosome.

2. *Mutation.* Informally, each chromosome mutates with probability $p_m > 0$. Formally, let $\{\varepsilon_t(i)\}_{i=1}^{M}$, with $\varepsilon_t(i) = (\varepsilon_{1t}(i),\ldots,\varepsilon_{nt}(i)) \in \mathbb{R}^n$, be M independent draws from a zero-mean, n-dimensional distribution. If chromosome i mutates then $\hat{\phi}_t(i) = \phi_t(i) + \varepsilon_t(i)$. Note that the mutations $\varepsilon_{kt}(i)$ may be correlated across genes, but not across chromosomes. We denote the mutation operator by μ.

3. *Tournament Selection.* Informally, the forecast performances of genes associated with randomly matched chromosomes are compared, resulting in a new chromosome made up of the superior genes. Forecast performance is determined by the average forecast error over the entire history. Formally, letting F denote the measure of PLM forecast performance, we set

$$F\left(\phi_{kt}(i), y^t\right) = -\frac{1}{t+1}\sum_{m=0}^{t}\left(y_{km} - \phi_{kt}(i)\right)^2.$$

   Here y^t is the history of y and $y_m = (y_{1m},\ldots,y_{nm})$ is the m-th period realization of y. Each period, M pairs of chromosomes are randomly selected *with replacement*. There are then two ways to proceed. Under *fine competition*, a new chromosome is created by selecting the genes of highest fitness. Thus if chromosomes i and j are paired and if $F\left(\phi_{kt}(i)\right) > F\left(\phi_{kt}(j)\right)$

then the k-th gene of the new chromosome is $\phi_{kt}(i)$. Under *coarse competition*, forecast performance is measured at the chromosome level by taking a weighted average of fitness across genes. In this case, the new chromosome is a copy of the chromosome with highest average fitness. We denote the tournament selection operator by τ.

Note that the period $t+1$ gene pool is determined by M independent draws from the evolved (via crossover and mutation) period t gene pool. In particular, the genes comprising a given agent's period $t+1$ chromosome depend on the period t gene pool, but not directly on the agent's period t chromosome. Further, under fine competition, the tournament selection operator produces a gene pool that does not depend on the affiliation of genes and chromosomes. It follows that, in linear models, fine competition makes crossover irrelevant.

Assuming, for concreteness, that we adopt the fine competition protocol, the model's dynamics under social learning are given by (3) and

$$\phi_{t+1} = \tau\left(\mu\left(\phi_t\right), y^t\right). \tag{4}$$

The stability question is the same as in the adaptive learning case: under what conditions do the $\phi_t(i)$ converge (in an appropriate sense) to zero? Under adaptive learning, and for appropriate gain sequences and model calibrations, almost sure convergence of the agents' PLMs to the REE obtains. Here, mutation precludes this possibility: the most that could be hoped for is convergence of the genes to an ergodic distribution centered at the REE.[6]

2.3 The E-stability Principle

Assume for the moment that agents are adaptive learners with homogeneous beliefs ϕ. In this case, the actual law of motion of the economy is

$$y_t = \beta\phi + v_t \equiv T(\phi) + v_t.$$

The T-map, $\phi \to \beta\phi$, can be interpreted as mapping the PLM to the ALM. At a fixed point of this map the PLM and ALM align, thus identifying the REE.

To assess whether we can expect adaptive learning agents to eventually coordinate their expectations on the REE, the learning literature advises analysis of the system of differential equations

$$\dot{\phi} = T(\phi) - \phi. \tag{5}$$

This simple dynamic system, sometimes referred to as *stylized adaptive learning*, dictates that beliefs evolve in the direction determined by a type of forecast error – the discrepancy between

[6] Arifovic, Bullard, and Kostyshyna (2013) assume that mutation variance decreases over time in order to allow for the possibility of a.s. convergence. In Section 4 we adopt their convention.

the PLM and the ALM, appropriately defined. Note that the rest point of this system characterizes the REE. If the rest point is Lyapunov stable then the REE is said to be *E-stable*, otherwise it is *E-unstable.* The *E-stability Principle* says that E-stable REE are learnable, i.e. stable under least-squares and related learning algorithms. The converse also holds – E-unstable REE are not learnable.

The E-stability principle is just that – a principle. Formally establishing that it holds for a given REE in a given model requires work – indeed, the relationship between (5) and adaptive learning dynamics is deep and technical, involving formalization of the ODE method. On the other hand, no important counter-examples to the principle are known, and the expedience it affords is considerable: simply compute the eigenvalues of DT at the rest point: if the real parts are less than one the REE is E-stable; if at least one eigenvalue has real part greater than one then the REE is E-unstable.

3 Learning in a univariate laboratory

In this section we examine learning dynamics, both adaptive and social, in the simplest possible environment: the univariate version of model of (1). In this world, chromosomes are genes, coarseness is fineness, and PLMs are real numbers. For notational simplicity, we let y^e_{it} be the PLM of agent i, that is, $E_t(i)y_{t+1} \equiv y^e_{it}$. Additionally we set $\bar{y}^e_t = M^{-1}\sum_{i=1}^M y^e_{it}$. The economy's dynamics are now given by $y_t = \beta\bar{y}^e_t + v_t$.

3.1 E-stability and adaptive learning

First we assess E-stability. Let y^e be the vector of agents' PLMs, with y^e_i being the PLM used by agent i. The actual law of motion can be written

$$y_t = \beta M^{-1} \cdot 1_M^\top \cdot y^e_t + v_t, \tag{6}$$

where $1_M \in \mathbb{R}^M$ is a column vector of ones. The induced T-map is

$$T(y^e) = \beta M^{-1}\left(1_M^\top \otimes 1_M\right) y^e.$$

Thus DT has an eigenvalue β associated to the eigenvector 1_M, and an $M-1$ dimensional kernel. It follows that the REE is E-stable if $\beta < 1$ and E-unstable if $\beta > 1$, as expected. In particular, by the E-stability Principle, $\beta > 1$ should preclude the REE as a long run outcome when the model is populated with adaptive learners.

3.1.1 Adaptive learning with decreasing gain

To examine this last point formally, assume agents update their beliefs using recursive least squares:

$$y^e_{it+1} = y^e_{it} + \gamma_t \left(y_t - y^e_{it} \right). \tag{7}$$

Equation (7) can be combined with $y_t = \beta \bar{y}^e_t + v_t$ to determine a dynamic system amenable to Ljung's theory of stochastic recursive algorithms. We have the following result, which is an instance of the E-stability principle in action:[7]

Theorem 1 (E-stability principle) *Assume agents update beliefs via RLS, and assume $\sum \gamma_t^2 < \infty$ and $\sum \gamma_t = \infty$.*

1. *If $\beta < 1$ then $y^e_{it} \to 0$ with probability one.*
2. *If $\beta > 1$ then $y^e_{it} \to 0$ with probability zero.*

3.1.2 Adaptive learning with constant gain

Social learning involves mutation – each period, each agent's beliefs are potentially subjected to perturbation. For comparison purposes, then, it is perhaps more natural to assume agents use a constant-gain algorithm: $\gamma \in (0,1)$, and usually taken to be small. For asymptotic results, the theory of stochastic recursive algorithms can be applied in this case as well, but to get a feel for transition dynamics, and thus to better understand how quickly we might expect instability to materialize when $\beta > 1$, it is convenient to assume $v_t \sim \mathcal{N}(0, \sigma_v^2)$ and to focus attention on the positive expectational-feedback case $\beta > 0$.

In the constant gain case, we may stack the recursions (7) and use the ALM to obtain a recursion for the expectations vector:

$$y^e_t = \left[(1-\gamma) I_M + \gamma M^{-1} \beta \left(1_M^\top \otimes 1_M \right) \right] y^e_{t-1} + \gamma 1_M v_t. \tag{8}$$

We see that y^e_t may be viewed as a VAR(1) process, which is stationary provided that the matrix corresponding to the term in the square brackets has eigenvalues inside the unit circle. A simple computation shows that the eigenvalues of this matrix comprise

$$\bar{\lambda} \equiv 1 + (\beta - 1)\gamma, \text{ and } M-1 \text{ copies of } 1-\gamma.$$

Tracking beliefs heterogeneity is also of interest. To this end, define the *beliefs dispersion* Δ^e_t at time t as

$$\Delta^e_t = M^{-1/2} \sqrt{\sum_i \left(y^e_{it} - \bar{y}^e_t \right)^2},$$

where $\bar{y}^e_t = M^{-1} 1_M^\top \cdot y^e_t$ is the cross-sectional mean of beliefs. We have the follow result:

[7] See Evans and Honkapohja (1996).

Proposition 1 *Given constant gain γ, the asymptotics of* (8) *are characterized as follows:*

1. *Beliefs dispersion converges to zero almost surely:* $\Delta_t^e \xrightarrow{a.s.} 0$.
2. *If* $\beta \in (0,1)$ *then*
   (a) *Beliefs y_{it}^e are asymptotically normally distributed around zero.*
   (b) *The state y_t converses weakly to a normal distribution:* $y_t \xrightarrow{\mathscr{D}} \mathscr{N}\left(0,(1+\xi)\sigma_v^2\right)$ *with* $\xi = \left(1-\bar{\lambda}^2\right)^{-1}\beta^2\gamma^2$, *where* $\bar{\lambda} \equiv 1+(\beta-1)\gamma$.
3. *If* $\beta > 1$ *then y_t is explosive.*

All proofs are in the Appendix. Item 2 can be made sharper. Since $\xi = \mathscr{O}(\gamma)$, we may conclude that for small γ the asymptotic distribution of the state under adaptive learning well-approximates the REE; further, conditional on initial beliefs, Ey_t converges to zero at rate $\bar{\lambda}$. See proof in the Appendix for details. Concerning item 3, by *explosive* we mean not uniformly bounded almost surely, and here also a sharper result is available: conditional on initial beliefs, $\|Ey_t^e\|$, and thus $|y_t|$, diverge at rate $\bar{\lambda}$.

Note that, by item 1, regardless of β's magnitude, the dynamics impart asymptotic homogeneity, reflecting that beliefs are adjusted in the direction of a common aggregate even if that aggregate is explosive. In fact, this result can be sharpened as well, though it is more naturally stated within the context of Proposition 2: see Corollary 1.

Proposition 1 provides details of the asymptotic behavior of the model under adaptive learning. Transitional behavior depends on initial beliefs. To broaden our assessment of transition dynamics we assume that the initial beliefs vector $y_0^e \in \mathbb{R}^M$ is obtained as a random sample of size M drawn from a normal distribution with zero mean and and finite variance $\sigma_0^2 > 0$. It is further assumed that the subsequent realizations of the shock v_t are independent of these draws.

Some of the results below are most naturally stated in terms of approximate distributions. We use the notation $x \dot{\sim} \mathscr{D}$ to indicate that the random variable x is approximately distributed as $\mathscr{D}$, where the nature of the approximation is context dependent. The following result characterizes the time t distributions of the state and of the dispersion of beliefs.

Proposition 2 *Let the initial condition $y_0^e \in \mathbb{R}^M$ be obtained as a random sample of size M drawn from* $\mathscr{N}\left(0,\sigma_0^2\right)$. *Then, provided M is sufficiently large,*

1. $y_t \sim \mathscr{N}\left(0,\Sigma_t^y\right)$ *where* $\Sigma_t^y = \left(\frac{\beta^2\sigma_0^2}{M} - \xi\sigma_v^2\right)\bar{\lambda}^{2t} + (1+\xi)\sigma_v^2$
2. $\Delta_t^e \dot{\sim} \mathscr{N}\left(\mu_t^{\Delta^e},\Sigma_t^{\Delta^e}\right)$ *where* $\mu_t^{\Delta^e} = \sqrt{\frac{2M-1}{2M}}\sqrt{\frac{M+1}{M}}(1-\gamma)^t\sigma_0 \approx (1-\gamma)^t\sigma_0$, *and*
$$\Sigma_t^{\Delta^e} = \frac{M+1}{2M^2}(1-\gamma)^{2t}\sigma_0^2 \approx \frac{1}{2M}(1-\gamma)^{2t}\sigma_0^2.$$

Item 2 provides for the following important corollary, referenced above:

Corollary 1 *Under the hypotheses of Proposition 2,* $\Delta_t^e \to 0$ *a.s. at a geometric rate regardless of whether the REE is stable under adaptive learning.*

Taken together, the above results suggest the following: if initial beliefs are drawn independently from a distribution centered at the mean of the REE then

1. The dispersion of beliefs should disappear: each agent's beliefs is expected to converge to mean beliefs.
2. If $\beta < 1$ then y_t should converge at rate $0 < \bar{\lambda} < 1$ to an ergodic distribution with mean equal to the mean of the REE.
3. If $\beta > 1$ then y_t should diverge at rate $\bar{\lambda} > 1$.
4. If $\beta \in (1, 1+\delta)$ for small $\delta > 0$ then the divergence of y_t should be very slow.

Figure 1 provides a simulation demonstrating some the findings of Propositions 1 and 2 in the unstable case, with heterogeneous initial beliefs symmetrically spaced about 0.0.[8] The blue line is the time path of mean beliefs; the time path of cross-sectional dispersion of beliefs is measured by plus/minus two standard deviations and is shown by the red lines. Note that expectations heterogeneity disappears quite quickly: the red curves overlay the blue curve by period 200. Also, the short simulation (left panel) gives the appearance of stability, whereas in the longer simulation (right panel) instability is evident.

3.2 E-stability and social learning

We turn now to social learning dynamics in our laboratory model. Because the model is univariate, crossover plays no role. We begin with a simulation, tuned to be roughly consistent with Arifovic, Bullard, and Kostyshyna (2013) – see Figure 2.[9] We consider the unstable case, $\beta = 1.01$, and assume that the economy has been at REE for 100 periods. This latter assumption affects the fitness assessment by providing an initial history at the REE.

Like Figure 1, in Figure 2 initial beliefs are symmetrically spaced about zero, and the blue line, which is very nearly zero and is covered by the red lines, is the time path of mean beliefs, with the red lines identifying two standard deviations. And also like Figure 1, the initial dispersion

[8] In this subsection we set the standard deviation of the exogenous shock to be small ($\sigma_v = 0.005$) in order to focus on instability induced by the learning mechanism. When constant gain is used, we set $\gamma = 0.05$.

[9] We set $p_m = 0.1$ and $M = 300$, though all results are qualitatively similar with $M = 30$, as in Arifovic, Bullard, and Kostyshyna (2013). We set the mutation standard deviation $\sigma_\varepsilon = 0.005$.

Figure 1. Adaptive learning: instability. Here the blue line is the time path of mean beliefs, and the red lines identify plus/minus two standard deviations in the cross-sectional dispersion of beliefs.

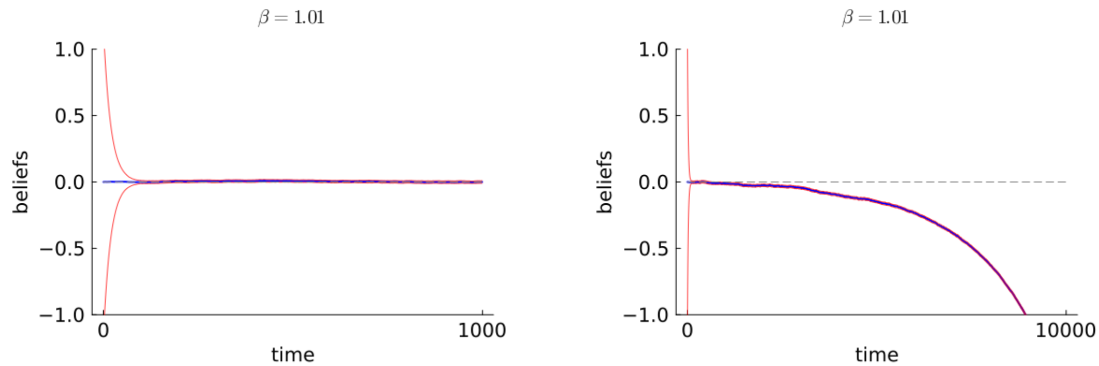

of beliefs is quickly eliminated. However, here we see that even after one million periods, there is no evidence of instability. This apparently conflicts with the expectational instability of the REE with $\beta = 1.01$ and is the matter this paper seeks to address. We refer to the phenomenon as *stable instability*.

Figure 2. Social learning: stable instability

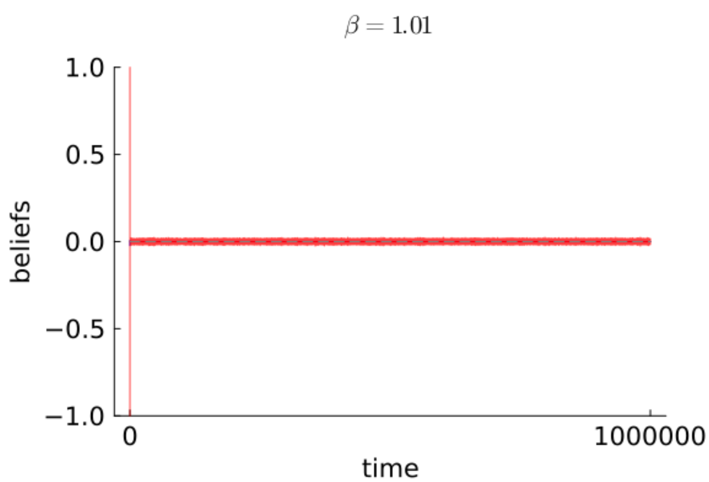

The stability analysis of Propositions 1 and 2 exploit the linear learning algorithm induced by constant gain RLS. More generally, stability results in the adaptive learning literature typically leverage the theory of stochastic recursive algorithms (SRAs). Interpreted through the lens developed by Evans and Honkapohja (2001), this theory is commonly applied to dynamic systems taking the following form:

$$\theta_t = \theta_{t-1} + \gamma_t H(\theta_{t-1}, x_t) + \gamma_t^2 \rho_t(\theta_{t-1}, x_t), \tag{9}$$

where $\theta_t \in \mathbb{R}^{n_1}$ is the estimator and $x_t \in \mathbb{R}^{n_2}$ is the state vector which is assumed to have a conditional distribution depending on x_{t-1} and θ_{t-1}. It is not obvious in general that social learning dynamics can be interpreted as a stochastic recursive algorithm, i.e. placed in the form (9); and this provides a possible explanation for Figure 2: maybe SL can be stable even when the REE is not E-stable.

In our simplified laboratory, and under some additional simplifying assumptions, it turns out that social learning *does* present as an SRA. Observe that

$$\begin{aligned} F\left(y_{it}^e, y^t\right) &= -(t+1)^{-1}\sum_{n=0}^{t}\left(y_{it}^e - y_n\right)^2 = -(t+1)^{-1}\sum_{n=0}^{t}\left(\left(y_{it}^e\right)^2 - 2y_{it}^e y_n + y_n^2\right) \\ &= -\left(\left(y_{it}^e\right)^2 - 2y_{it}^e (t+1)^{-1}\sum_{n=0}^{t} y_n + (t+1)^{-1}\sum_{n=0}^{t} y_n^2\right) \\ &= 2y_{it}^e \bar{y}_t - \left(y_{it}^e\right)^2 + \text{ terms common across agents.} \end{aligned} \tag{10}$$

It follows that we can use $2y_{it}^e \bar{y}_t - \left(y_{it}^e\right)^2$ as our fitness measure. Abusing notation somewhat, we write $\tau(y_t^e, \bar{y}_t)$ as the tournament selection operator based on this new fitness measure.

Putting it all together results in the following dynamic system:

$$\bar{y}_t = \bar{y}_{t-1} + (t+1)^{-1}\left(\beta M^{-1}\langle 1_M, y_t^e\rangle + v_t - \bar{y}_{t-1}\right) \tag{11}$$

$$y_{t+1}^e = \tau(\mu(y_t^e), \bar{y}_t). \tag{12}$$

By interpreting $\bar{y}_t$ as the estimator and y_t^e as the state vector, this system takes the form an SRA with Markovian state dynamics, and is thus amenable to the ODE method.

To compute the algorithm's mean dynamics, we must have sufficient understanding of the asymptotic behavior of (12) for fixed $\bar{y}$. The following lemma sets the stage by saying that, for fixed $\bar{y}$, the state dynamics (12) is stable ergodic, with asymptotic mean beliefs equal to $\bar{y}$.

Lemma 1 *Assume the mutation density is uniform with compact, connected support, and that each agent's PLM mutates each period ($p_m = 1$). Fix $\bar{y}$ and let y_t^e evolve according to* (12)*. Then there is a distribution $\nu(\bar{y})$ over beliefs such that $y_t \to y^e \sim \nu(\bar{y})$ weakly. Furthermore, $E_{\nu(\bar{y})} y_i^e = \bar{y}$.*

The proof is the Appendix.[10] We can now study the stability properties of the SRA (11) - (12). The mean dynamics are given by

$$\frac{d\bar{y}}{dt} = h(\bar{y}) \equiv \int \left(\beta M^{-1}\langle 1_M, y^e\rangle - \bar{y}\right)\nu(\bar{y})(dy^e). \tag{13}$$

[10]Simulations according with the assumptions of the theorem, i.e. uniform draws for mutation and $p_m = 1$, yield the same stable instability as is evidenced in Figure 2 – indeed the associated figure is qualitatively indistinguishable: see Appendix.

By Lemma 1 we see that $h(\bar{y}) = (\beta - 1)\bar{y}$. We conclude that, under our simplifying assumptions, the stability properties of the mean dynamics associated with social learning align with those implied by E-stability.[11]

Lemma 1 is cold comfort: after all, Keynes' admonition that we're all dead in the long run holds in spades in Figure 2. It also raises an interesting question: What explains stable instability?

3.3 Explaining stable instability

Stable instability arises in a univariate model, so cross-over is not the culprit. Also, heterogeneity of expectations cannot be central – see Figure 1 panel b; and Figure 2 demonstrates that widely distributed (heterogeneous) initial beliefs do not induce instability, strongly suggesting that mutation cannot be the primary driver of stable instability. Thus we view tournament play as our prime suspect, and proceed with its interrogation.

The metric for tournament play is the fitness measure F, which, by equation (10), may be interpreted as implementing a tendency to move the estimator towards the sample mean of past data. Importantly, this is in contrast to our implementations of adaptive learning, which tend to move the estimator towards the most recent realization of the data. This observation suggests a two-pronged experimental approach to explaining stable instability. First, we modify adaptive learning by implementing an algorithm that moves the estimator towards the mean of past data, and see if we can induce stable instability. And second, we modify social learning by using a fitness measure that rewards proximity to the most recent realization of the data, rather than to the mean of past data, and see if we get instability as in the usual adaptive learning case.

3.3.1 Modified adaptive learning

The updating model (7) is modified in two ways: first, an idiosyncratic shock is introduced to simulate mutation; and second, agents update their estimates based on the mean of all past data rather than only on the newest data point. This latter modification simulates tournament play, as will be explained in more detail below.

The model dynamics are now given by (6) and

$$\bar{y}_t = \bar{y}_{t-1} + \hat{\gamma}_t (y_t - \bar{y}_{t-1}) \tag{14}$$

$$y^e_{it} = y^e_{it-1} + \gamma \left(\bar{y}_t - y^e_{it-1} + \varepsilon_{it}\right). \tag{15}$$

Equation (14) provides the recursive estimation of the sample mean, associated with gain $\hat{\gamma}_t$. Equation (15), which comprises the principal distinction between this section and the previous, says

[11] To formally connect the stability properties of the mean dynamics (13) with those of the SRA requires demonstrating a host of conditions are met: see assumptions M.1 – M.5 on pages 155 - 157 of Evans and Honkapohja (2001). We leave this analysis to the motivated reader.

that agents move a mutation – i.e. a perturbation given by ε_{it} – of their forecasts in the direction of the sample mean. The ε_{it} are assumed mean zero, independently distributed across space and time, and to have small support.

Equation (15) is meant to capture, at least in spirit, tournament behavior. In a tournament, mutated forecasts are matched, and the best forecast is selected to perpetuate. In REE, $y_t = v_t$, and thus the data represent iid draws, whence the best forecast is the sample mean: $\hat{\gamma}_t = t^{-1}$. Equation (15) says that agents move their mutated estimate in the direction of the best forecast.

The E-stability principle holds for the modified model. In particular,

Proposition 3 *Assume $\sum \hat{\gamma}_t = \infty$ and $\sum \hat{\gamma}_t^2 < \infty$. Under the dynamics* (6), (14)*, and* (15)*,*

1. *If $0 < \beta < 1$ then $\bar{y}_t$ converges to zero with probability one.*
2. *If $\beta > 1$ then $\bar{y}_t$ converges to zero with probability zero.*

Also, under constant gain, $\hat{\gamma}_t = \gamma$, the analogs of Propositions 1 & 2 can be developed. Stack the dynamic system as before to get

$$\begin{pmatrix} y_t^e \\ \bar{y}_t \end{pmatrix} = B \cdot \begin{pmatrix} y_{t-1}^e \\ \bar{y}_{t-1} \end{pmatrix} + \begin{pmatrix} \gamma I_M & \gamma^2 1_M \\ 0 & \gamma \end{pmatrix} \begin{pmatrix} \varepsilon_t \\ v_t \end{pmatrix}, \tag{16}$$

where

$$B = \begin{pmatrix} (1-\gamma) I_M + \frac{\beta \gamma^2}{M} \left(1_M^\top \otimes 1_M \right) & \gamma(1-\gamma) I_M \\ \frac{\beta \gamma}{M} 1_M^\top & 1-\gamma \end{pmatrix}.$$

The eigenvalues of B are $M-1$ copies of $1-\gamma$ and $\lambda^{\pm}$, where $\lambda^- < 1-\gamma < \lambda^+$. The following proposition summarizing the results.

Proposition 4 *Assume $\hat{\gamma}_t = \gamma$, and that initial values of y_0^e are drawn as in Proposition 2. The dynamics implied by the system (16) satisfy the following properties:*

1. *If $\beta \in (0,1)$ then y_t and y_{it}^e are asymptotically normally distributed around zero, and Ey_t converges to zero at rate $0 < \lambda^+ < 1$ where $\lambda^+ > \underbrace{1-\gamma+\beta\gamma}_{\bar{\lambda}}$.*
2. *If $\beta > 1$ then $E|y_t|$ diverges at rate $1 < \lambda^+ < \underbrace{1-\gamma+\beta\gamma}_{\bar{\lambda}}$.*
3. *If mutation is shut down, i.e. $\varepsilon_{it} = 0$, then beliefs dispersion converges to zero almost surely. Furthermore, $\Delta_t^e \dot{\sim} \mathcal{N}\left(\mu_t^{\Delta^e}, \Sigma_t^{\Delta^e}\right)$ where $\mu_t^{\Delta^e} \approx (1-\gamma)^t \sigma_0$ and $\Sigma_t^{\Delta^e} \approx \frac{1}{2M}(1-\gamma)^{2t}\sigma_0^2$.*
4. *If $\beta \in (0,1)$ and $\sigma_\varepsilon > 0$ then, asymptotically, $\Delta_t^e \sim \mathcal{N}\left(\sqrt{\frac{2M-1}{2M}}\sqrt{\frac{M+1}{M}}\frac{\sigma_\varepsilon}{\sqrt{2\gamma(1-\gamma)}}, \frac{M+1}{4M^2}\frac{\sigma_\varepsilon^2}{2\gamma(1-\gamma)}\right)$.*

Thus, in the constant gain case, stability, i.e. convergence to a stationary distribution, continues to turn on the size of β relative to one. Also, the dispersion of beliefs does not depend on v_t, and is eliminated at exactly the same rate as in the previous section.

Items 1 and 2 of the proposition requires some further discussion. Recall, from the discussion following Proposition 1 indicates that the rate of convergence/divergence of y_t is governed by $\bar{\lambda} = 1 - \gamma + \beta\gamma$. Items 1 and 2 of Proposition 4 thus say that the system (16) converges more slowly than the usual adaptive learning dynamics when $\beta < 1$, and, more importantly, *diverges more slowly* when $\beta > 1$.

Figure 3 provides clear evidence that the modified adaptive learning mechanisms exhibits stable instability of a form similar to social learning. Here $\hat{\gamma}_t = t^{-1}$, and the figure should be compared to Figure 1, which provides a simulation under the usual adaptive learning dynamics, and to Figure 2, which show social learning dynamics. Note that, consistent with Figure 2, in Figure 3 there is no evidence of instability even after ten million periods.[12]

Figure 3. Modified adaptive learning: stable instability

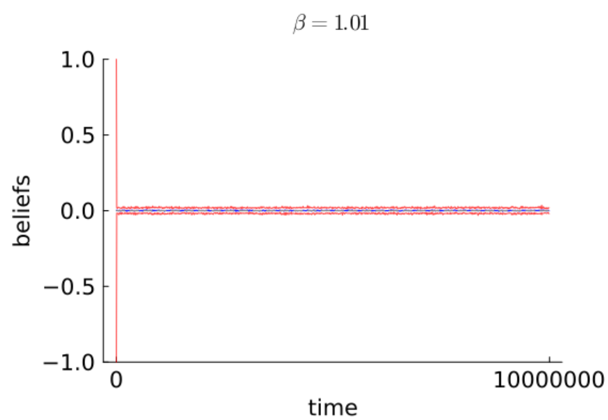

3.3.2 Modified social learning

In Section 3.3.1, we modified the standard adaptive learning algorithm by dictating the estimator be adjusted in accordance with its discrepancy from the sample mean of all past data, rather than from the most recent realization of the data. This modification, in effect, respects the fitness measure used in social learning. Here we work the opposite direction by modifying the fitness measure used in SL to accord with the standard adaptive learning algorithm. Specifically, we set

$$F\left(y_{it}^e, y^t\right) = -\left(y_{it}^e - y_t\right)^2 = 2y_{it}^e y_t - \left(y_{it}^e\right)^2 + \text{ terms common across agents }, \tag{17}$$

[12]Under adaptive learning modified to include mutation, we scale the σ_ε down by a factor of five, which, in the stable case, results in an ergodic distribution of beliefs with roughly the same variance as that obtained under social learning.

which should be compared with equation (10). In effect, beliefs nearest the most recent realization of the endogenous variable are deemed fittest.

Figure 4 provides a simulation of the economy under modified social learning. This figure should be compared with Figures 1 and 2. In particular, this figure demonstrates that modified social learning has stability properties more closely aligned with the usual adaptive learning mechanism – instability after 10000 periods is clearly evidenced.[13]

A caveat merits comment. While the time path over the first 1000 periods always suggests stability, there is considerable variation in the behavior of over the next 9000 periods, with some simulations indicating continued stability and others becoming unstable more quickly that is shown in the right panel of the figure. A statistical analysis over multiple simulations would be revealing.

Figure 4. Modified social learning: instability

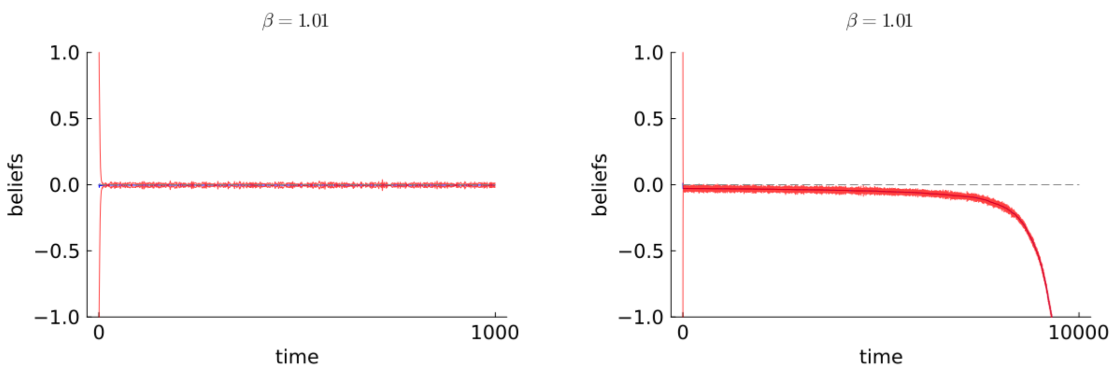

3.3.3 Stable instability: discussion

Stable instability is a vaguely defined, transient phenomenon that does not lend itself to precise explanation. It is not unique to social learning, and can arise under any learning mechanism if the underlying model has expectational feedback sufficiently near unity.

In the macroeconomics literature, social learning has been found to induce stable instability in instances where adaptive learning would not. Our two-pronged analysis here suggests that these findings may reflect the chosen specification of the fitness measure, rather than the social learning dynamic more broadly. In the next section, we test this hypothesis by considering social learning – standard and modified – in the same New Keynesian environment studied by Arifovic, Bullard, and Kostyshyna (2013).

[13] Arifovic, Grimaud, Salle, and Vermandel (2020) employ a social learning algorithm with a fitness measure that discounts past forecast errors. Our fitness measure, (17), relies solely on the most recent forecast error since past forecast errors are already reflected in current forecasts. Arifovic, Salle, and Truong (2023) also use a fitness measure that conditions on contemporaneous errors to better fit experimental data derived from a New Keynesian environment.

4 Social learning in a New Keynesian model

In this section we examine social learning and stable instability in a benchmark New Keynesian model. Our model is chosen to align with Arifovic, Bullard, and Kostyshyna (2013) (ABK). In deviation form, it is given by

$$\begin{array}{ll} \text{IS:} & y_t = E_t y_{t+1} - \sigma^{-1}(i_t - E_t \pi_{t+1}) \\ \text{AS:} & \pi_t = \delta E_t \pi_{t+1} + \kappa y_t \\ \text{PR:} & i_t = \varphi_\pi \pi_t + \varphi_y y_t \end{array}$$

Here y_t is log deviation of output from its flexible price level, π_t is log deviation of the inflation factor from the inflation target π^*, and i_t is the log deviation of the interest rate from its target i^*.

This model can be placed in the form (1) by eliminating the contemporaneously determined interest rate and solving for output and inflation in terms of expectations. We have

$$\begin{pmatrix} y_t \\ \pi_t \end{pmatrix} = \beta E_t \begin{pmatrix} y_{t+1} \\ \pi_{t+1} \end{pmatrix}, \text{ where} \tag{18}$$

$$\beta = \frac{1}{\sigma + \varphi_y + \kappa\varphi_\pi} \begin{pmatrix} \sigma & 1 - \delta\varphi_\pi \\ \kappa\sigma & \kappa + \delta(\sigma + \varphi_y) \end{pmatrix}.$$

We denote by ξ the eigenvalue of β with largest magnitude – it captures the expectational feedback in the model, and so pins down the model's determinacy and E-stability properties.

Because the model is purely forward-looking and absent exogenous shocks, the MSV solution is $y_t = \pi_t = 0$. Bullard and Mitra (2002) demonstrate that this MSV is the unique REE and it is stable under adaptive learning exactly when monetary policy is *active*, i.e. when the following version of the Taylor principle is satisfied:

$$\kappa(\varphi_\pi - 1) + (1 - \delta)\varphi_y > 0. \tag{19}$$

The main finding of ABK is that the MSV solution is stable under social learning regardless of whether the Taylor principle is satisfied.

4.1 Stable instability in the NK model

We begin by reproducing the findings of ABK. Like them, we modify the mutation operator to allow for a diminishing variance in the perturbation: using their notation, the period t mutation $\hat{\varepsilon}_t$ is taken as

$$\hat{\varepsilon}_t = \big(1 - (\text{decrease}) \times t/T\big)\varepsilon_t,$$

where T is the length of the simulation and decrease $\in (0,1)$ measures the rate at which the variance diminishes over time. We also adopt their quarterly calibration of the model: $\delta = 0.99$, $\sigma = 0.157$, and $\kappa = 0.024$.

Figure 5 replicates the passive policy figure of ABK: the simulation length is 1000 periods, with $M = 30$ agents, and policy parameters set as $\varphi_\pi = \varphi_y = 0.5$.[14] These parameters result in a value of -0.007 for the LHS of equation (19) – the Taylor Principle is not satisfied: the model's steady state is indeterminate and E-unstable.[15] The feedback parameter's value is $\xi = 1.01$, which is the same feedback used in our ad hoc univariate model.

Figure 5. Stable instability under SL with passive policy

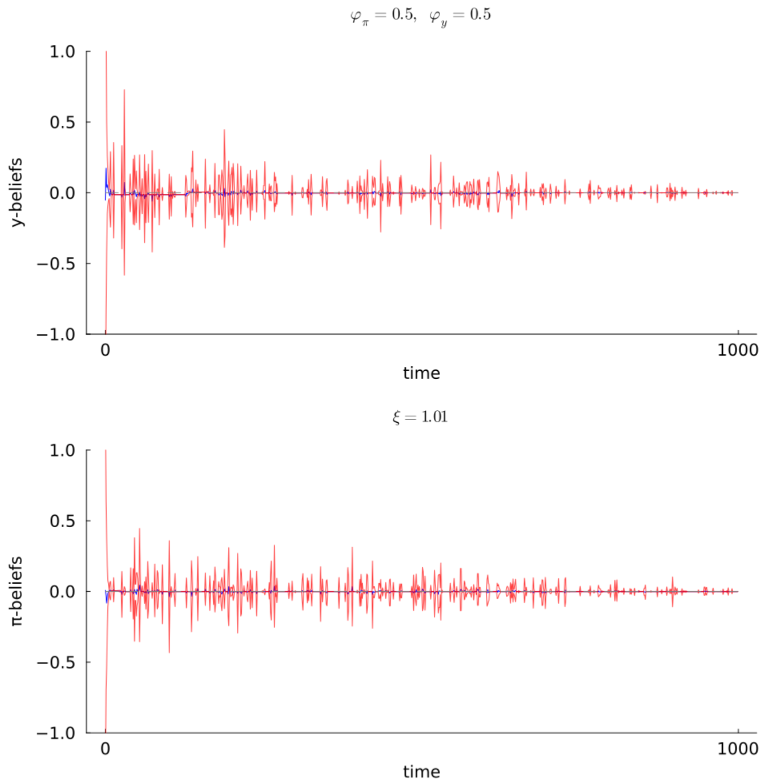

In fact, numerical analysis shows there is no sign of divergence even after 1,000,000 periods. This is a remarkable finding. The monetary policy literature has long argued the imperative that the Taylor Principle be satisfied. Admittedly, the reasoning behind the argument has evolved somewhat. It (seems to have) started with the compelling intuition that policy makers should raise rates more than one-to-one when faced with increased inflation, "lean against the wind" policy that results in rising real rates, thus cooling the economy. The argument favoring the Taylor principle became more rigorous when it was discovered that passive policy implies indeterminacy, and thus raises the possibility that self-fulfilling "sunspot" equilibria may introduce increased volatility. Finally, Bullard and Mitra (2002) put the nail in the coffin by showing that the Taylor principle was

[14] All results are robust to using $M = 300$ agents, as in Arifovic, Schmitt-Grohe, and Uribe (2018).

[15] Some examinations of monetary policy include interactions with fiscal policy, which can affect the determinacy and E-stability properties of the model's steady state. We follow ABK and (implicitly) assume Ricardian fiscal policy.

necessary for the MSV solution to be stable under adaptive learning. Figure 5 counters 20 years of policy advice: if agents expectations evolve via a social learning dynamic then passive policy can induce stable instability.

Figure 6 provides a simulation under modified social learning. As expected, instability emerges over time horizons that are consistent with instability under adaptive learning. It merits observing that the initialization of beliefs, which are drawn randomly using the mutation distribution, play a significant role in the ensuing dynamics, and this leads to considerable variation in outcomes. While instability always emerges, initializations that are very nearly centered at pre-policy-change MSV levels lead to longer periods of transient stability. For this reason, increasing the number of agents tends to increase the length of time before instability emerges.

Figure 6. Stable instability under modified SL with passive policy

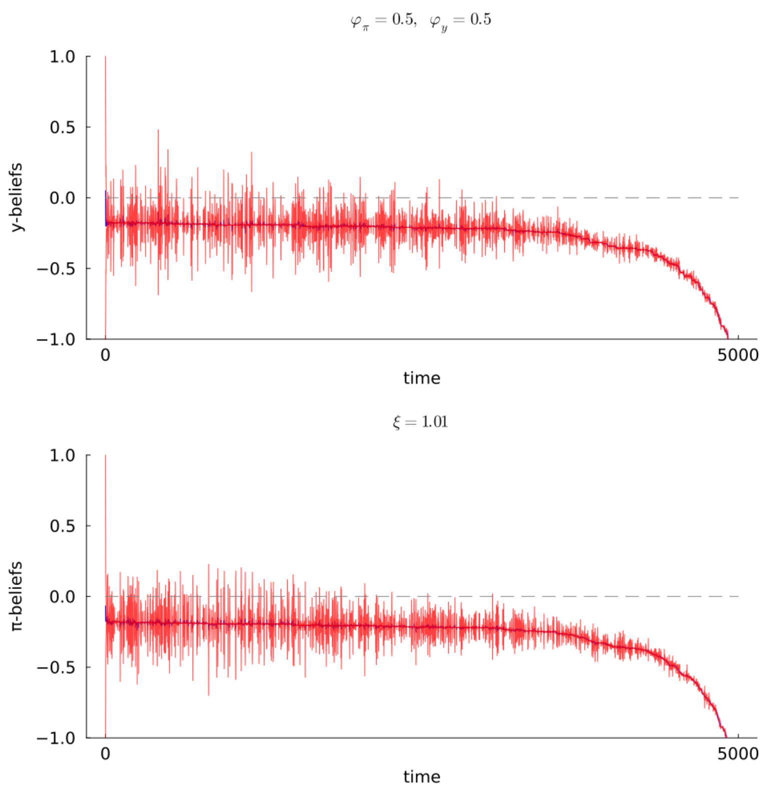

4.2 Policy changes and social learning

In an economy populated with social learning agents, the central bank can use passive policy to keep the economy near the MSV solution. This raises a natural question: can passive policy be used to implement, say, an announced permanent change in the inflation target? The results so far do not suggest an answer to this question: under passive policy social learners do not deviate from the MSV, but can they learn a new one?

To examine this question, we consider the following policy experiment that is perhaps a propos to the times: we assume the inflation target has been at 5% for 100 periods, and then the central bank announces an unexpected, permanent decrease in the target to 0%. Standard policy prescription would say to implement this change by using an instrument rule that satisfies the Taylor principle. This implementation is simulated in Figure 7. We note that realized output and inflation very closely track mean output and inflation beliefs, so we refrain from plotting them.

Figure 7. Lowering the inflation target under SL with active policy

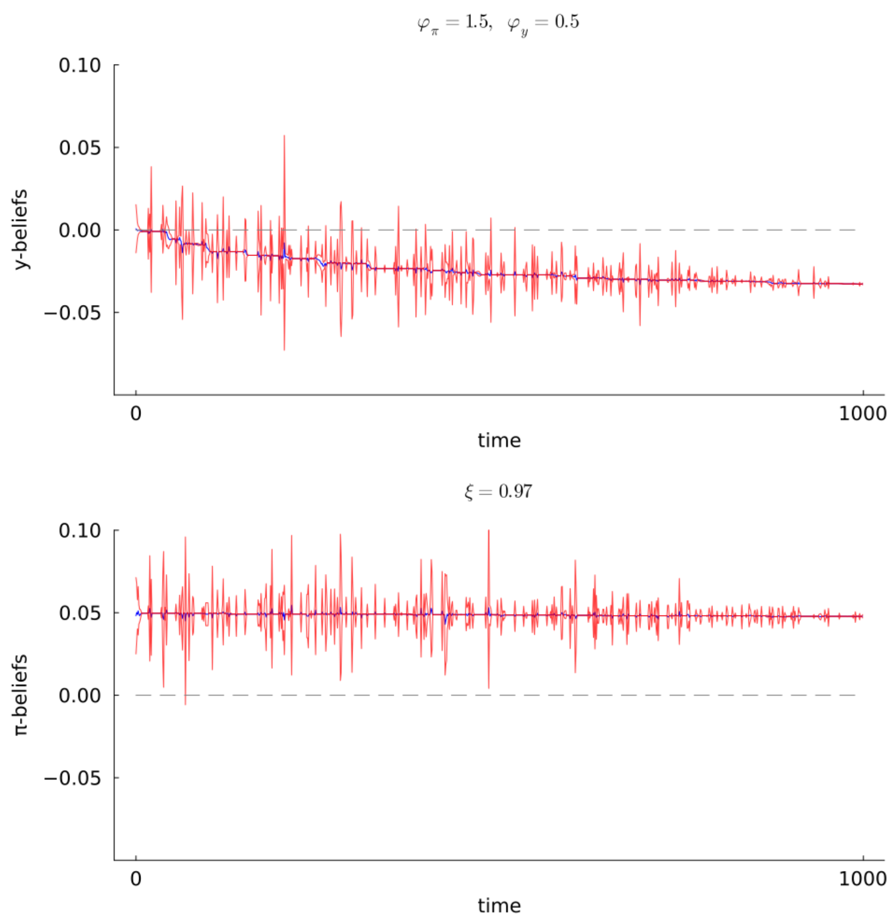

Here we have chosen an inflation coefficient in the instrument rule of $\varphi_\pi = 1.5$, which leads to a feedback magnitude of $\xi = 0.97 < 1$, guaranteeing determinacy and stability under adaptive learning.

We see that the policy does not result in a decline in inflation, at least over the first 1000 periods; however output falls substantially. When the central bank lowers the inflation target, the current relative inflation level is suddenly high, thus the instrument rule requires that the nominal rate is raised. Since the Taylor Principle is satisfied, this rise causes an increase in the real rate, and hence a reduction in output. The failure of inflation to fall in part reflects the significant weight placed on past inflation history, but it also reflects the slope of the Phillips curve: current inflation puts a weight of 0.99 on expected inflation and only 0.024 on current output. Running the simulation for 1,000,000 periods does not change the result.

Figure 8 provides a simulation under passive policy. In this case, the outcome is the same for inflation, but the opposite for output. The reasoning is exactly the same as before: the fall in the inflation target leads to a current relative inflation level that is suddenly high; however, this time policy is passive, so the instrument rule raises the nominal rate less than one-to-one, resulting in a fall in the real rate and a rise in output. The stability of inflation beliefs follows just as above. Running the simulation for 1,000,000 periods does not change the result.

Figure 8. Lowering the inflation target under SL with passive policy

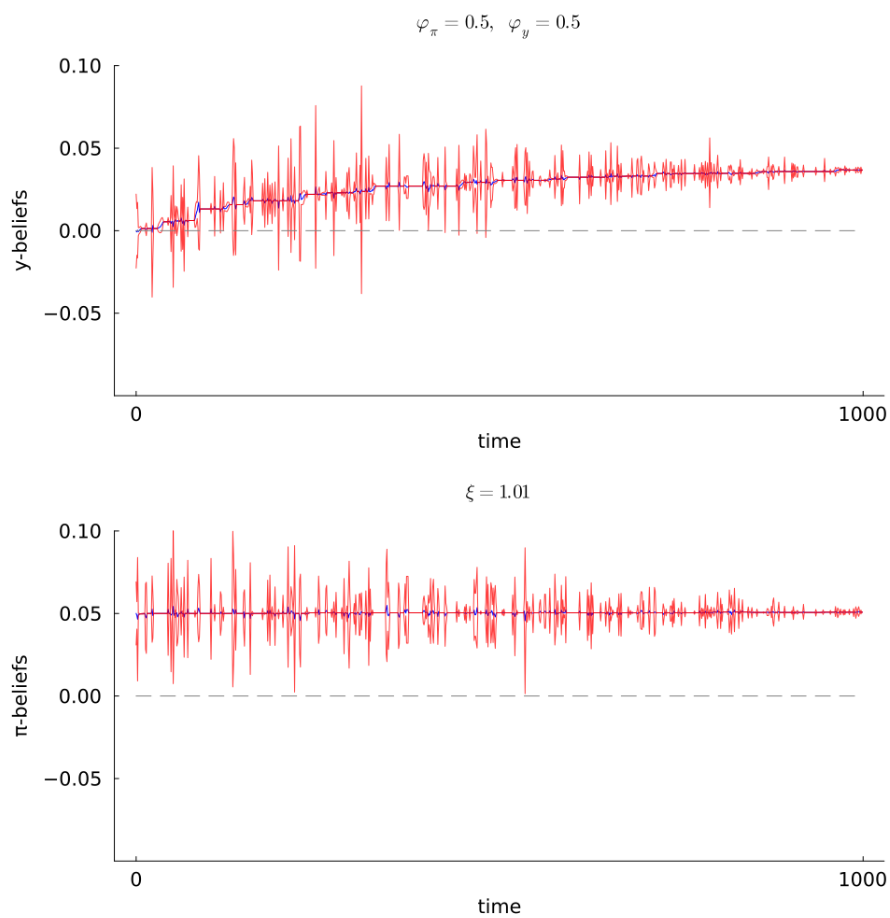

The failure of even active policy to move the economy to the new MSV solution is in part a reflection of the stable instability induced by social learning. This observation suggests that active policy coupled with modified social learning might lead the economy to the new inflation steady state. Figure 9 provides a simulation of exactly this experiment, and the results are as expected. The lowering of the target leads to a reduction in output, just as before, but now inflation beliefs adjust downward, raising output over time toward its steady state.

Figure 9. Lowering the inflation target under modified SL with active policy

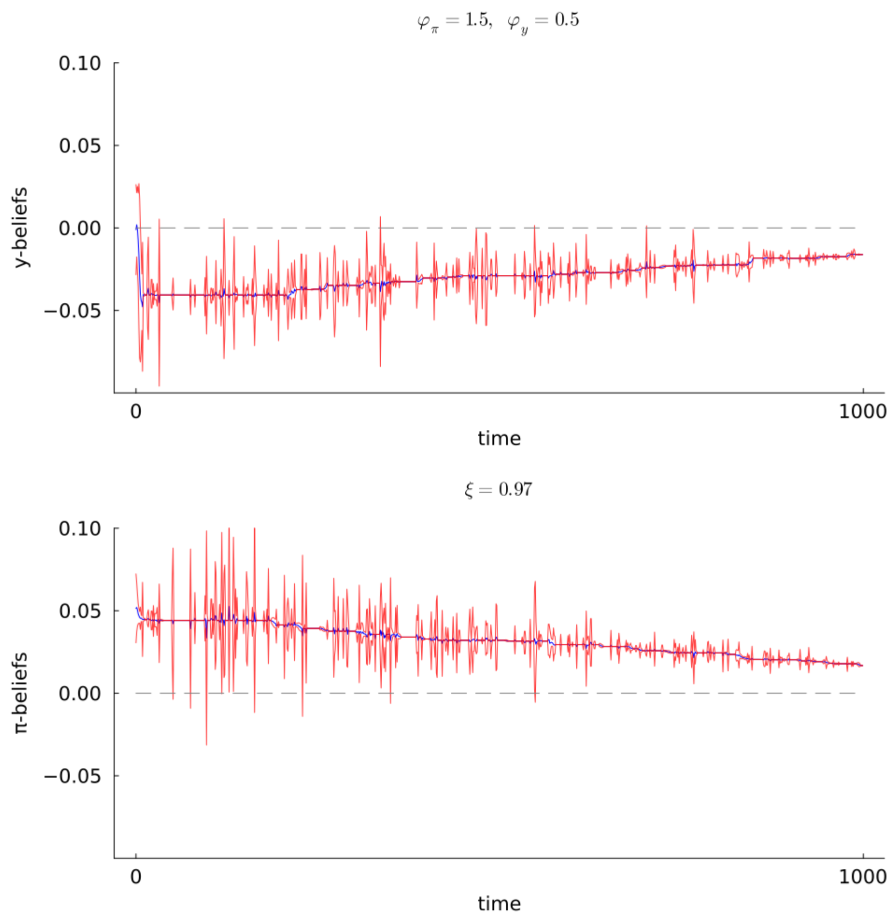

The very different responses, under SL, AL and modified SL, to an unanticipated permanent reduction in the inflation target, are striking and forcefully illustrate the role of the fitness measure. However, alternative scenarios can be imagined, and quite different results would likely arise from an unanticipated persistent but *temporary* shock. AL allows for this possibility by using a constant gain, tuned to capture the trade-off between tracking structural change and filtering noise, whereas SL emphasizes long-term fitness measures. Bullard (2023) argues that the latter could be an advantage in adjusting to large but temporary shocks such as the Great Recession and the Pandemic recession and subsequent inflation. The argument is that precisely because SL has genes strongly adapted to the normal steady state, when a large but temporary disruption ends, the economy will more quickly leave the disruption behind.

These considerations suggest that both SL and AL approaches have advantages (which may or may not be aligned at the individual and aggregate levels) in dealing with unforeseen large temporary shocks. While this paper has focused on differences in fitness measures, it may be equally important that the forecasting and decision rules in use include significant heterogeneity, possibly reflecting very different past states of the economy, to ensure their alertness to changing conditions.

5 Conclusion

Social Learning is an attractive and flexible learning model that has been used to study a wide range of issues in macroeconomics, and other areas, including exchange rates, growth transitions and finance. A major contribution of this paper has been to show that, at least with some specific technical assumptions, SL can be set up as a stochastic recursive algorithm. This makes it amenable to the ODE method, and in particular enables us to obtain mean dynamics in line with the E-stability approach. Using this procedure we showed that asymptotic SL results are in accordance with the expectational stability principle. This is important because E-stability is usually simple to compute and has provided a reliable way to assess asymptotic behavior of learning algorithms in economic models. A second contribution of this paper has been to identify the differing fitness criteria explicitly or implicitly used in AL and SL as the source for the "stable instability" results under passive monetary policy that have been observed under SL.

The analysis presented here can be extended a number of natural ways, and along three particular dimensions. First, the model under study could be enriched to include additional explanatory variables, i.e. serially correlated observable shocks and/or lagged endogenous variables. Second, within the models studied in the paper alternative forecast rules could be considered, including behavioral rules along the lines found in experimental works: see Hommes (2013) discussion and many details. Finally, alternative fitness measures could be considered, including measures that condition on discounted lagged forecast errors, as in Arifovic, Grimaud, Salle, and Vermandel (2020).[16]

Our focus on stable instability in the New Keynesian model reflects its importance to monetary policy, as well as our intention to address and explain the apparently conflicting results found in earlier studies, Arifovic, Bullard, and Kostyshyna (2013) in particular. It would be natural, and of considerable interest, to examine asset-pricing applications, especially since the shorter time-scale found in financial markets makes the possibility of stable instability all the more relevant.[17]

[16] Using this fitness measure, the cited paper obtains the intriguing numerical result in a non-linear NK model that the basin of attraction of the (locally determinate) targeted steady state is larger under SL than it is under AL.

[17] For example, Branch and Evans (2011) and Williams (2023) consider stock-market models in which prices stay for long periods near the fundamentals price, but occasionally follow bubble paths which endogenously pop and return to the fundamentals price.

We think SL and AL should be viewed as complementary approaches. AL can include (and has included) heterogeneous forecasting rules. SL could in principle encode least-squares forecasting rules as genes. A major issue for both approaches going forward is how to deal with structural change. Because structural change can be permanent, temporary or recurring, there may be advantages in embedding both current-oriented and long-term-oriented fitness measures in learning models.

References

ARIFOVIC, J. (1994): "Genetic Algorithm Learning and the Cobweb Model," *Journal of Economic Dynamics and Control*, 18, 3–28.

——— (1995): "Genetic Algorithms and Inflationary Economies," *Journal of Monetary Economics*, 36, 219–243.

——— (1996): "The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies," *Journal of Political Economy*, 104, 510–541.

ARIFOVIC, J., J. BULLARD, AND J. DUFFY (1997): "The Transition from Stagnation to Growth: An Adaptive Learning Approach," *Journal of Economic Growth*, 2, 185–209.

ARIFOVIC, J., J. BULLARD, AND O. KOSTYSHYNA (2013): "Social Learning and Monetary Policy Rules," *Economic Journal*, 123, 38–76.

ARIFOVIC, J., A. GRIMAUD, I. SALLE, AND G. VERMANDEL (2020): "Social learning and monetary policy at the effective lower bound," Discussion paper.

ARIFOVIC, J., I. SALLE, AND H. TRUONG (2023): "History-Dependent Monetary Regimes: A Lab Experiment and a HENK Model," Tinbergen Institute Discussion Paper TI 2023-028/VI, Amsterdam and Rotterdam.

ARIFOVIC, J., S. SCHMITT-GROHE, AND M. URIBE (2018): "Learning to Live in a Liquidity Trap," *Journal of Economic Dynamics and Control*, 89, 120–136.

BENVENISTE, A., M. METIVIER, AND P. PRIOURET (1990): *Adaptive Algorithms and Stochastic Approximations*. Springer-Verlag, Berlin.

BRANCH, W. (2006): "Restricted Perceptions Equilibria and Learning in Macroeconomics," in Colander (2006), pp. 135–160.

BRANCH, W. A., AND G. W. EVANS (2011): "Learning about Risk and Return: A Simple Model of Bubbles and Crashes," *American Economic Journal: Macroeconomics*, 3, 159–191.

BRAY, M., AND N. SAVIN (1986): "Rational Expectations Equilibria, Learning, and Model Specification," *Econometrica*, 54, 1129–1160.

BULLARD, J. (2023): "Social Learning for the Masses," Presentation at the Computational & Experimental Economics Workshop, Simon Fraser University, Feb. 4, 2023.

BULLARD, J., AND K. MITRA (2002): "Learning About Monetary Policy Rules," *Journal of Monetary Economics*, 49, 1105–1129.

COLANDER, D. (2006): *Post Walrasian Macroeconomics*. Cambridge, Cambridge, U.K.

EVANS, G. W. (1985): “Expectational Stability and the Multiple Equilibria Problem in Linear Rational Expectations Models,” *The Quarterly Journal of Economics*, 100, 1217–1233.

——— (1989): “The Fragility of Sunspots and Bubbles,” *Journal of Monetary Economics*, 23, 297–317.

EVANS, G. W., AND S. HONKAPOHJA (1996): “Least Squares Learning with Heterogeneous Expectations,” *Economics Letters*, 53, 197–201.

——— (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, New Jersey.

EVANS, G. W., S. HONKAPOHJA, AND R. MARIMON (2001): “Convergence in Monetary Inflation Models with Heterogeneous Learning Rules,” *Macroeconomic Dynamics*, 5, 1–31.

EVANS, G. W., AND B. MCGOUGH (2020): “Adaptive Learning and Macroeconomics,” *Oxford Research Encyclopedia of Economics and Finance*, https://doi.org/10.1093/acrefore/9780190625979.013.508.

HOMMES, C. (2013): *Behavioral Rationality and Heterogeneous Expectations in Complex Economics Systems*. Cambridge University Press, Cambridge, UK.

KUSHNER, H., AND D. CLARK (1978): *Stochastic Approximation Methods for Constrained and Unconstrained Systems*. Springer-Verlag, Berlin.

LJUNG, L. (1977): “Analysis of Recursive Stochastic Algorithms,” *IEEE Transactions on Automatic Control*, 22, 551–575.

MARCET, A., AND T. J. SARGENT (1989): “Convergence of Least-Squares Learning Mechanisms in Self-Referential Linear Stochastic Models,” *Journal of Economic Theory*, 48, 337–368.

ROBBINS, H., AND S. MONRO (1951): “A Stochastic Approximation Method,” *Annals of Mathematical Statistics*, 22, 400–407.

SARGENT, T. J. (1993): *Bounded Rationality in Macroeconomics*. Oxford University Press, Oxford.

——— (2008): “Evolution and Intelligent Design,” *American Economic Review*, 98, 5–37.

WILLIAMS, N. (2023): “A Learning Model of Financial Instability,” Working paper.

WOODFORD, M. (1990): “Learning to Believe in Sunspots,” *Econometrica*, 58, 277–307.

Appendix

Proof of Proposition 1 . Begin by diagonalizing the system (8):

$$(1-\gamma)I_M + \gamma M^{-1}\beta\left(1_M^\top \otimes 1_M\right) = S\Lambda S^{-1},$$

with Λ denoting the diagonal matrix of eigenvalues ordered such that

$$\Lambda_{MM} = 1-\gamma(1-\beta) \equiv \bar{\lambda}. \tag{20}$$

With help from a computer algebra system, it is straightforward to show that we may write

$$S^{-1} = \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \\ 0 & 0 & 0 & \cdots & 0 \end{pmatrix} + \frac{1}{M}\begin{pmatrix} -1 & -1 & -1 & \cdots & -1 \\ -1 & -1 & -1 & \cdots & -1 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ -1 & -1 & -1 & \cdots & -1 \\ 1 & 1 & 1 & \cdots & 1 \end{pmatrix}.$$

Noting

$$S^{-1}y_t^e = (y_{2t}^e - \bar{y}_t^e, \cdots, y_{Mt}^e - \bar{y}_t^e, \bar{y}^e)',$$

we change coordinates to get

$$\begin{aligned} y_{it}^e - \bar{y}_t^e &= (1-\gamma)\left(y_{it-1}^e - \bar{y}_{t-1}^e\right) \text{ for } i = 1,\cdots,M & (21)\\ \bar{y}_t^e &= \bar{\lambda}\bar{y}_{t-1}^e + \gamma v_t. & (22) \end{aligned}$$

Now observe that the dynamics of $y_{it}^e - \bar{y}_t^e$ for $i=1$ may be inferred from the dynamics for $i = 2,\cdots,M$ by noting that $\sum_i (y_{it}^e - \bar{y}_t^e) = 0$. Thus, it follows from equation (21) that, conditional on y_0^e, the dynamics of beliefs dispersion are deterministic, and they converge to zero regardless of β, establishing item 1. Items 2 and 3 follow by noting, from equation (20), that $|\bar{\lambda}| < 1 \Leftrightarrow \beta < 1$. ■

Proof of Proposition 2. Equations (21) and (22) continue to determine the dynamics, thus

$$\bar{y}_t^e = \bar{\lambda}^t \bar{y}_0 + \gamma \sum_{n=1}^{t} \bar{\lambda}^{t-n} v_n.$$

By our assumptions on the determination of initial conditions we have $\bar{y}_0 \sim \mathcal{N}\left(0, M^{-1}\sigma_0^2\right)$. It follows that $\bar{y}_t^e$ is normally distributed, with mean zero and variance

$$\Sigma_t^{\bar{y}} = \frac{\sigma_0^2}{M}\bar{\lambda}^{2t} + \gamma^2\left(\frac{1-\bar{\lambda}^{2t}}{1-\bar{\lambda}^2}\right)\sigma_v^2,$$

which, combined with $y_t = \beta\bar{y}_{t-1}^e + v_t$, establishes item 1.

Turning to item 2, first note that provided M is sufficiently large, we may take y^e_{i0} and $\bar{y}^e_0$ as approximately independent. It follows that

$$y^e_{i0} - \bar{y}^e_0 \dot{\sim} \mathscr{N}\left(0, \left(\frac{M+1}{M}\right)\sigma_0^2\right).$$

Iterating (21) we have

$$y^e_{it} - \bar{y}^e_t = (1-\gamma)^t \left(y^e_{i0} - \bar{y}^e_0\right) \dot{\sim} \mathscr{N}\left(0, \left(\frac{M+1}{M}\right)(1-\gamma)^{2t}\sigma_0^2\right).$$

Now let $\psi_t^2 = M^{-1}(M+1)(1-\gamma)^{2t}\sigma_0^2$ so that $\psi_t^{-1}(y^e_{it} - \bar{y}^e_t)$ is standard normal. It follows that $M\psi_t^{-2}(\Delta^e_t)^2 \dot{\sim} \chi^2(M)$. Item 2 of proposition 2 is established using the Fisher approximation, which says that for large M we have $x \sim \chi^2(M)$ implies $\sqrt{2x} \dot{\sim} \mathscr{N}\left(\sqrt{2M-1}, 1\right)$. ■

Proof of Lemma 1. Fix $\bar{y}_t$ and y^e_t. We need to understand the distribution $\sigma(\mu(y^e_t), \bar{y}_t)$ in the special case that all agents hold the same forecast: $y^e_{it} = y^e_{jt} \equiv y^e$, where we drop the time subscript for the remainder of this step to thin notation. The distribution of interest, $\sigma(\mu(y^e), \bar{y})$, characterizes the outcomes of the following procedure: two independent draws, ε_1 and ε_2, from the mutation perturbation identify two candidate forecasts $x_1 = y^e + \varepsilon_1$ and $x_2 = y^e + \varepsilon_2$; the candidate forecast nearest to $\bar{y}$ is selected.

Let f be the density of random mutation $y^e \to y^e + \varepsilon$, with F the associated distribution. To determine the density function associated to the distribution $\sigma(\mu(y^e), \bar{y})$, let $x \in \mathbb{R}$ represent an arbitrary forecast, and ask the following question: what is the probability that a mutation of y^e is *worse* that x? First suppose $x < \bar{y}$. Then a random draw $x' = y^e + \varepsilon'$ is worse, i.e. farther from $\bar{y}$, if it is less that x or if it is greater than $2\bar{y} - x$. Symmetric reasoning addresses the case $\bar{y} < x$. We conclude that the density function of interest, $d\sigma$,is given by

$$d\sigma(x) = \begin{cases} 2\left(F(x) + 1 - F(2\bar{y} - x)\right) f(x) & x \leq \bar{y} \\ 2\left(1 - F(x) + F(2\bar{y} - x)\right) f(x) & x \geq \bar{y} \end{cases} \tag{23}$$

Figure 10 provides graphical intuition for this formula. The approximate probabilities of selecting two distinct points, $x < 0 < \tilde{x}$, are considered – the former in blue and the latter in red. The solid blue lines identify the set of all points (x, x') and (x', x) in (x_1, x_2)-space for which x' is worse than x. The dashed blue rectangles approximate the events that one of these points is the outcome of the random draw $(y^e + \varepsilon_1, y^e + \varepsilon_2)$ of candidate forecast models.

For fixed $\bar{y}$ and y^e_t, let

$$y^e_{t+n+1} = \tau\left(\mu\left(y^e_{t+n}\right), \bar{y}\right). \tag{24}$$

We adjoin a projection facility to this dynamic to ensure y^e_{t+n} remains in a compact ball *cube* centered at $\bar{y}$. The radius of this ball can be taken as arbitrarily large – we only need compactness for the argument to hold.

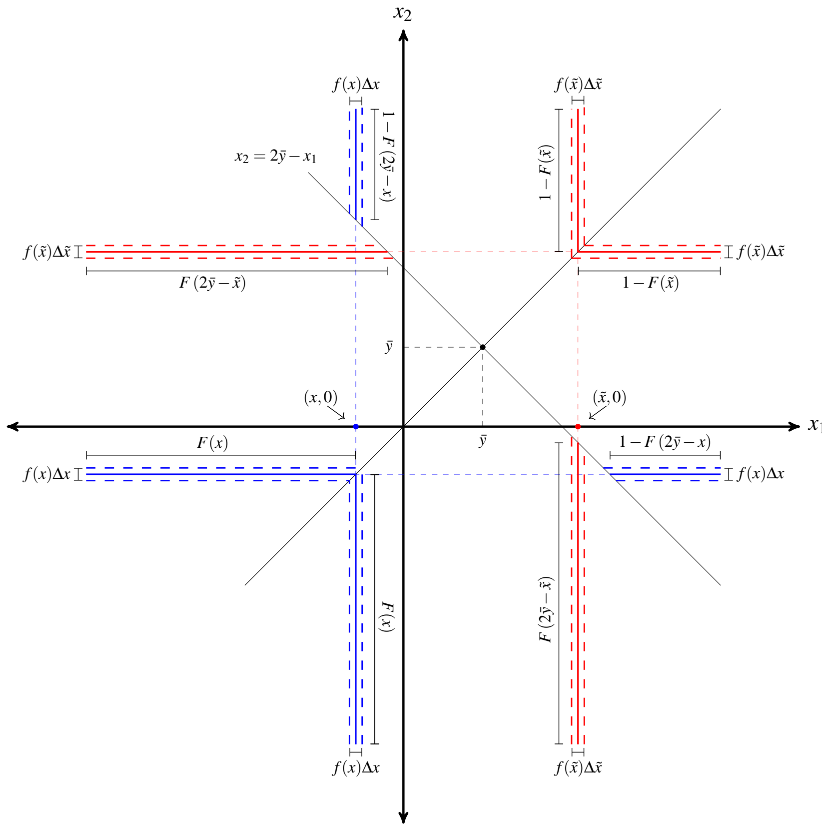

Figure 10. Selection density

Our first goal is to establish that (24) is stable ergodic, i.e. y^e_{t+n} convergences weakly to a random vector distributed as $\Gamma_{\bar{y}}$. Denote by $\mathscr{B}$ the Borel subsets of B, and by $Q_{\bar{y}}(\cdot,\cdot)$ the Markov transition characterized by (24):

$$Q_{\bar{y}}(y^e, A) = \text{prob}\big(\tau\left(\mu\left(y^e\right), \bar{y}\right) \in A\big), \text{ for all } y^e \in B,\ A \in \mathscr{B}. \tag{25}$$

By Doeblin minorization, to demonstrate stable ergodicity it suffices to find $N \in \mathbb{N}$ and a measure λ on B, such that

$$Q^N_{\bar{y}}\left(y^e, A\right) \geq \lambda\left(A\right), \text{ for all } y^e \in B,\ A \in \mathscr{B}, \tag{26}$$

where Q^N is the N-step iteration of Q. To this end, expanding the space B if necessary, embed a lattice in B with N_1 uniformly distributed nodes such that neighboring nodes differ only in one coordinate, with the difference set at $\Delta/3$; and cover B in N_1 open cubes $\{B_i\}_{i=1}^{N_1}$ of side-length $\Delta/2$ centered on the nodes: see Figure 11 for illustration. By construction (from the overlapping nature of the cover and from the mutation distribution properties, i.e. uniform with compact, connected support), it follows that for any $n > N_1$, there is a $\delta(n) > 0$ such that the probability of moving from B_i to B_j in n steps is greater than $\delta(n)$. In particular, $n > N_1 \implies Q^n_{\bar{y}}(x, B_i) > \delta(n)$ for all $x \in B$ and all B_i.

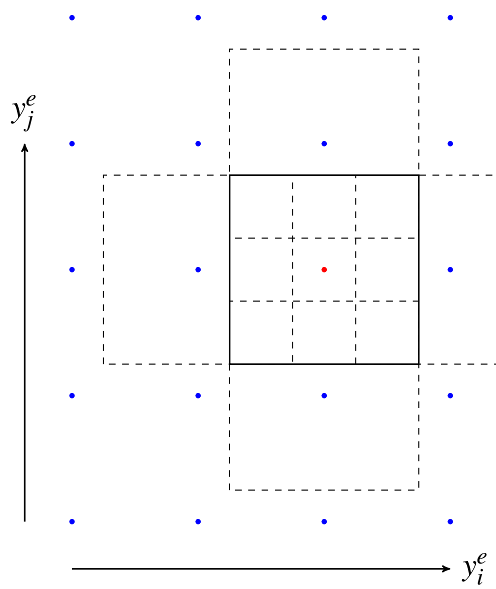

Figure 11. Open cover figure

Now let $\hat{\lambda}$ be Lesbesgue measure on B. Note that it suffices to show (26) holds for open cubes, so let $A \in \mathscr{B}$ be any open cube, with, say center at $\hat{x} \in B$. Then $\hat{x} \in B_i$ for some element of the cover. Let y^e be in the closure of B_i and note that $Q_{\bar{y}}(y^e, dx)$ is the conditional density of our Markov process. Because B_i has side-length $\Delta/2$, we know that this density is strictly positive on the closure of $A \cap B_i$. Let $Q_{\bar{y}}^{\min}(y^e)$ be its minimum value. Since $Q_{\bar{y}}^{\min}(\cdot)$ is continuous in y^e, we may let $Q_{\bar{y}}^{\min}$ be the minimum over the closure of B_i. We conclude that

$$y^e \in B_i \implies Q(y^e, A \cap B_i) \geq Q_{\bar{y}}^{\min} \cdot \hat{\lambda}(A \cap B_i).$$

Since the center of A is in B_i, it follows that there is a $\hat{\delta} > 0$ so that $\hat{\lambda}(A \cap B_i) \geq \hat{\delta}\hat{\lambda}(A)$. Putting this all together, we can take $N = N_1 + 1$ and $\lambda = \delta(N)\hat{\delta}\hat{\lambda}$, and conclude that (26) holds.

We have established weak convergence of y_t^e to a random vector distributed as $\Gamma_{\bar{y}}$. Denoting its mean by $(\bar{y}_1^e, \ldots, \bar{y}_M^e)$, we have, from the key insight, that $\bar{y}_i^e = \bar{y}_j^e \equiv \bar{y}^e$. We claim that $\bar{y}^e = \bar{y}$. To see this, first note that equation (24) implies $\bar{y}^e = E\sigma(\mu(\bar{y}^e), \bar{y})$, and that the density of $\sigma(\mu(\bar{y}^e), \bar{y})$ is given by (23), with f a uniform distribution, centered at $\bar{y}^e$, with support $\Delta > 0$. Direct computation shows[18]

$$E\sigma(\mu(\bar{y}^e), \bar{y}) = \begin{cases} \bar{y}^e + \frac{\Delta}{6} & \bar{y}^e \leq \bar{y} - \Delta/2 \\ \frac{4(\bar{y}-\bar{y}^e)^3}{3\Delta^2} - \frac{2(\bar{y}-\bar{y}^e)^2}{\Delta} + \bar{y} & \bar{y} - \Delta/2 < \bar{y}^e \leq \bar{y} \\ \frac{4(\bar{y}-\bar{y}^e)^3}{3\Delta^2} + \frac{2(\bar{y}-\bar{y}^e)2}{\Delta} + \bar{y} & \bar{y} < \bar{y}^e \leq \bar{y} + \Delta/2 \\ \bar{y}^e - \frac{\Delta}{6} & \bar{y}^e > \bar{y} + \Delta/2 \end{cases} \tag{27}$$

[18]This is where we are using $p_m = 1$.

We see that $\bar{y}^e = E\sigma\left(\mu\left(\bar{y}^e\right), \bar{y}\right)$ implies $\bar{y} - \Delta/2 < \bar{y}^e < \bar{y} + \Delta/2$, so that

$$\bar{y}^e = \frac{4(\bar{y} - \bar{y}^e)^3}{3\Delta^2} \pm \frac{2(\bar{y} - \bar{y}^e)^2}{\Delta} + \bar{y}. \tag{28}$$

Divide each side of (28) by Δ, set the equation to zero, and let $\rho = \Delta^{-1}(\bar{y} - \bar{y}^e)$, to get

$$\rho\left(\frac{4}{3}\rho^2 \pm 2\rho + 1\right) = 0. \tag{29}$$

Since roots of the quadratics are necessarily complex, we conclude that $\rho = 0$, i.e. $\bar{y}^e = \bar{y}$, which establishes our claim. ■

Proof of Proposition 3. We may write the dynamic as

$$\bar{y}_t = \bar{y}_{t-1} + \hat{\gamma}_t\left(\frac{\beta}{M}1_M^\top y_{t-1}^e - \bar{y}_{t-1} + v_t\right) \tag{30}$$

$$y_t^e = (1-\gamma)y_{t-1}^e + \gamma 1_M \bar{y}_t + \gamma\varepsilon_t \tag{31}$$

Now let $x_t = (y_{t-1}^e, v_t)$ and $\xi_t = (1, \varepsilon_{t-1}, v_t)$. Then, for fixed $\bar{y}$, the state dynamics (31) becomes

$$x_{t+1} = ((1-\gamma)I_M \oplus 0)x_t + \begin{pmatrix} \gamma 1_M \bar{y} & \gamma I_M & 0 \\ 0 & 0 & 1 \end{pmatrix} \xi_{t+1}.$$

It follows that the system (30)-(31) has the appropriate SRA form for application of the ODE method. Thus let

$$H(\bar{y}, x_t) = \frac{\beta}{M}1_M^\top y_{t-1}^e - \bar{y} + v_t.$$

For fixed $\bar{y}$, we may compute, using 31, that $Ey_t^e = \bar{y}$, so that

$$h(\bar{y}) = \lim_{t\to\infty} EH(\bar{y}, x_t) = (\beta - 1)\bar{y}.$$

The ode ${}^{d\bar{y}}\!/_{d\tau} = h(\bar{y})$ is Lyapunov stable at $\bar{y} = 0$ provided $\beta < 1$, and is not Lyapunov stable if $\beta > 1$. ■

Sketch of Proof of Proposition 4. The arguments here are similar to those for Propositions 1 and 2, and so they are sketched. Let $\psi = \left(4(1-\gamma) + \beta\gamma^2\right)^{-1/2}$. Using Mathematica, we may determine that the eigenvalues of B (see (16)) are $M-1$ copies of $1-\gamma$ and

$$\lambda^{\pm} = 1 - \gamma + \frac{\beta\gamma^2}{2} \pm \frac{\gamma}{2}\frac{\sqrt{\beta}}{\psi},$$

from which the stated stationarity conditions and properties of $\lambda^{\pm}$ follow. Continuing as above, we may write $B = S\Lambda S^{-1}$ with

$$S^{-1} = \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 0 & \cdots & 0 & 0 \end{pmatrix} + \frac{1}{N} \begin{pmatrix} -1 & -1 & -1 & \cdots & -1 & 0 \\ -1 & -1 & -1 & \cdots & -1 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ -1 & -1 & -1 & \cdots & -1 & 0 \\ 0 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 0 & \cdots & 0 & 0 \end{pmatrix}$$

$$+ \frac{\psi}{M} \begin{pmatrix} 0 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 0 & 0 \\ -1 & -1 & -1 & \cdots & -1 & \frac{M}{2\psi}\left(\beta^{-1/2} + \gamma\psi\right) \\ 1 & 1 & 1 & \cdots & 1 & \frac{M}{2\psi}\left(\beta^{-1/2} - \gamma\psi\right) \end{pmatrix}.$$

The remaining results follow from this decomposition. ■