

IMPACTS OF PRIORS ON CONVERGENCE AND ESCAPES FROM NASH INFLATION

THOMAS J. SARGENT AND NOAH WILLIAMS

ABSTRACT. Recent papers have analyzed how adaptive agents may converge to and escape from self-confirming equilibria. All of these papers have imputed to agents a particular prior about drifting coefficients. In the context of a model of monetary policy, this paper analyzes dynamics that govern both convergence and escape under a more general class of priors for the government. We characterize how the shape of the prior influences the dynamics in important ways. There are priors for which the E-stability condition is not enough to assure local convergence to a self-confirming equilibrium. Our analysis also tracks down the source of differences in the sustainability of Ramsey inflation encountered in the analyses of Sims (1988) and Chung (1990), on the one hand, and Cho, Williams, and Sargent (2002), on the other.

Key Words: Self-confirming equilibrium, mean dynamics, escape route, large deviation, natural rate of unemployment, adaptation, priors.

1. INTRODUCTION

A widely used justification for assuming rational expectations is that with a long enough historical data record, statistical learning will equate objective and subjective probability distributions, at least on events that are observed along an equilibrium path. The only possible survivors of a learning process whereby decision makers recalibrate and dispose of poorly fitting models are elements of a particular *subset* of rational expectations equilibria, called *self-confirming equilibria* (SCE).¹ In a self-confirming equilibrium, agents use models that cannot be improved in light of the observations that occur within the equilibrium. But decision makers' models can be wrong off an equilibrium path. If a large agent (like the government in a macro model) has a wrong model, then that agent's misreading of occurrences off an equilibrium path can influence equilibrium outcomes.

Sargent (1999) defined a self-confirming equilibrium for a version of the Phillips curve model of Kydland and Prescott (1977). The self-confirming equilibrium inflation outcome happens to equal the *Nash inflation level*, namely, Kydland and Prescott's time-consistent outcome.² The self-confirming equilibrium outcome is supported by a misspecified government model (a non-expectational Phillips curve) that in equilibrium is observationally equivalent with the true data generating model (an expectational Phillips curve model), but that disagrees with the true model off the equilibrium path. The government's model can be described by a vector γ of regression coefficients.

Date: April 30, 2003.

We thank Jess Benhabib, Jim Bullard, Lars Peter Hansen, Seppo Honkapohja, and Tao Zha for helpful comments. Pierre-Olivier Weill provided research assistance.

¹This insight pervades Fudenberg and Levine (1998).

²In Kydland and Prescott's model, the government knows the correct model.

A self-confirming equilibrium can emerge as a limit point of an adaptive model in which the government recursively refits its possibly misspecified model as data accrue. In this adaptive model, the government designs the regression coefficients γ to estimate a linear relationship whose coefficients α the government believes drift over time according to

$$(1.1) \quad \alpha_n = \alpha_{n-1} + \Lambda_n,$$

where Λ_n is an i.i.d. vector process with mean zero and covariance matrix V_n . The government acts as a Bayesian who maintains (1.1) as part of its prior. The regression coefficients γ are updated over time and represent the government's estimate of the true underlying, time-varying parameters α . Sargent (1999) and Cho, Williams, and Sargent (2002) established conditions under which such an adaptive model would converge to a self-confirming equilibrium. Cho, Williams, and Sargent (2002) did this for specifications of V_n of two types, leading to belief updating rules known as "constant gain" (in which V_n is constant over time) and "decreasing gain" (in which V_n shrinks over time), respectively. Convergence is governed by the so-called "mean dynamics," an ordinary differential equation that summarizes the relentless action of least squares to bring the government's model into line with the historical data. For $V_n = 0$, the mean dynamics predict that the adaptive model converges to a SCE.

But Cho, Williams, and Sargent (2002) and Sargent (1999) discovered that a second component, the so-called "escape dynamics," recurrently pushes the government's model *away* from a self-confirming equilibrium and toward beliefs that support the low-inflation *Ramsey outcome*. The escape dynamics are governed by another ODE. Cho, Williams, and Sargent (2002) analyzed the escape dynamics for their particular adaptive learning algorithms and used them to explain simulations that had been prepared by Sargent (1999). With decreasing gain estimators like least squares, the escape dynamics gradually expire, but they endure with constant gain algorithms. Williams (2002) provides a complete characterization of the escape dynamics, determining their frequency and identifying the most likely path that beliefs follow when they escape from a self-confirming equilibrium.

This paper studies convergence and escape dynamics under alternative possible specifications for the coefficient innovation covariance matrix V_n in equation (1.1). We focus on cases in which this matrix is time-invariant and study the effects of different specifications of it. The prior covariance matrix shapes the directions of variation in the parameters along which the government most suspects movements. We show that different prior beliefs can lead to substantively different outcomes. Interestingly, we find that the convergence of beliefs is no longer governed by the "E-stability" conditions stressed by Evans and Honkapohja (2001), but instead depends on prior beliefs about parameter drift. In particular, some belief specifications lead to instability and cycles in a model that is "E-stable".³ Further, we show how the specification of beliefs affects the speeds of convergence and escape, and the directions in which escapes occur.

We also address unanswered questions about discrepancies between the simulations of closely related models of Sargent (1999) and Cho, Williams, and Sargent (2002), on the one hand, and Sims (1988) and Chung (1990), on the other. While the former two studies find

³Bullard (1994) analyzes the existence of similar cycles due to learning in an overlapping generations model. Honkapohja and Mitra (2002) also extend the E-stability conditions in their analysis of learning with heterogeneous agents.

that escapes from Nash inflation are always followed by subsequent movement back toward the self-confirming Nash equilibrium as the mean dynamics reassert themselves, the latter two studies exhibit simulations that start from the self-confirming Nash outcome, escape to the low inflation Ramsey outcome, but then apparently stay there indefinitely. What accounts for the failure of the mean dynamics to reassert themselves in the simulations of Sims and Chung?⁴ This is a substantively important question because by perpetuating inflation-unemployment outcomes superior to the Nash outcome, the simulations of Sims and Chung suggest that, despite its specification error, the government can permanently learn a version of the natural rate hypothesis that is good enough to inspire good actions.

To answer these questions, this paper extends the analysis of Cho, Williams, and Sargent (2002) by allowing V_n in (1.1) to take more general forms. We analyze recursive algorithms in which the government uses a Kalman filter, rather than the more restrictive recursive least squares imposed in Cho, Williams, and Sargent (2002). That allows us to get closer to the specifications of Sims (1988) and Chung (1990). By letting us match their specifications of priors, we have the tools needed to track down how Sims and Chung arrested the mean dynamics after an escape.

The remainder of this paper is organized as follows. Section 2 describes a model of inflation and a self-confirming equilibrium. Section 3 specifies the government’s view about coefficient drift, while section 4 shows the relationship between some adaptive algorithms and the Kalman filter. Section 5 gives examples of some priors for which an adaptive algorithm cycles in spite of E-stability. Sections 6 and 7 analyze the impacts of priors on the speed of convergence to a self-confirming equilibrium and on the nature of escapes from it. Section 8 explains unusual features of the simulations of Sims (1988) and Chung (1990), while section 9 concludes.

2. SETUP

This section defines a self-confirming equilibrium and then describes the basic model that we’ll be using.

2.1. Self-confirming equilibrium. A self-confirming equilibrium has these components:

- A true data generating model.
- An erroneous model that describes the government’s beliefs and guides its decisions both on and off an equilibrium path.
- A self-confirmation requirement that makes the government’s model and the true data generating mechanism agree along the equilibrium path.

In our example, we take the true data generating process to be a version of the natural rate hypothesis under rational expectations, the government’s model to be a regression equation that portrays an empirical Phillips curve, and the self-confirmation requirement to be that the government’s regression equation matches that implied by the true data generating mechanism. We now fill in the details.

⁴Their papers do not pose this question, perhaps because they offer no formal analysis of convergence or escape.

2.2. The truth. Let $W_n = [W_{1n} \ W_{2n}]'$ be a sequence of i.i.d. Gaussian vectors with mean zero and covariance I . Let U_n be the unemployment rate and π_n the inflation rate at time n . Let x_n be the systematic part of the inflation rate, which can be set by the government, and let \hat{x}_n be the public's expectation of π_n . The economy is described by the following version of a model of Kydland and Prescott (1977):

$$(2.2) \quad U_n = u - (\pi_n - \hat{x}_n) + \sigma_1 W_{1n}, \quad u > 0$$

$$(2.3) \quad \pi_n = x_n + \sigma_2 W_{2n}$$

$$(2.4) \quad \hat{x}_n = x_n.$$

Equation (2.2) is a natural rate Phillips curve, where u is the natural rate of unemployment and $\pi_n - \hat{x}_n$ is the public's error in forecasting inflation. Equation (2.3) equates realized inflation to the government's setting of target inflation x_n plus a random disturbance. Equation (2.4) imposes rational expectations for the public.

2.3. The government's beliefs and its control problem. The government has a model that is observationally equivalent with the truth along an equilibrium stochastic process, but nevertheless wrong because of what it assumes about behavior *off* the equilibrium process. Define the vector

$$(2.5) \quad X_{n-1} = [U_{n-1} \ U_{n-2} \ \pi_{n-1} \ \pi_{n-2} \ 1]'$$

The government's model of the economy is a linear Phillips curve with parameters $\gamma = [\gamma'_{-1} \ \gamma_1]'$:

$$(2.6) \quad U_n = \gamma_1 \pi_n + \gamma'_{-1} X_{n-1} + \eta_n$$

where the government treats η_n as a mean zero, serially uncorrelated random term that is beyond its control and that is orthogonal to X_{n-1} .⁵ Although our paper is about alternative ways to restrict γ , we temporarily regard it as arbitrary. The government's decision rule (2.8) solves the problem:

$$(2.7) \quad \min_{\{x_n\}} \hat{E} \sum_{n=0}^{\infty} \delta^n (U_n^2 + \pi_n^2)$$

where \hat{E} denotes the expectations operator induced by model (2.6) and the minimization is subject to (2.6) and (2.3). The solution to this problem is a government best response of the form

$$(2.8) \quad x_n = h(\gamma)' X_{n-1}.$$

We call problem (2.7) the *Phelps problem*. Sargent (1999) and CWS (2002) identify three salient alternative beliefs and the outcomes associated with them:

⁵For expository purposes, we shall also consider the simpler model in which the government estimates a static regression of unemployment on inflation and a constant (i.e. $X_{n-1} = 1$). We call this the *static model*. Since there is no temporal dependence in (2.2),(2.3), all of the temporal dependence in the model comes through the government's beliefs. For the static model, the government's decision rule can be calculated explicitly, allowing some of our characterizations to be sharper.

- Belief 1. If $\gamma_1 = -1, \gamma_{-1} = [0 \ 0 \ 0 \ 0 \ u]$, then the Phelps problem tells the government to set $x_n = u$ for all n . This is the Nash outcome of Sargent (1999), i.e., the time-consistent outcome of Kydland and Prescott (1977).
- Belief 2. If $\gamma_1 = 0, \gamma_{-1} = [0 \ 0 \ 0 \ 0 \ u^*]$ for any u^* , the government sets $x_n = 0$ for all n . This is the Ramsey outcome, i.e., the optimal time-inconsistent outcome of Kydland and Prescott (1977).
- Belief 3. If the coefficients on current and lagged π_n 's sum to zero, then as $\delta \rightarrow 1$ from below, the Phelps problem eventually sends x_n arbitrarily close to 0.

Under the actual probability distribution generated by (2.2), (2.3), (2.4), the value of the government's objective function (2.7) is larger under the outcome $x_n = 0$ than under $x_n = u > 0$. Under Belief 1, the government perceives a trade-off between inflation and unemployment and to exploit that trade-off sets inflation above zero. Under Belief 2, the government perceives no trade-off, sets inflation at zero, and accepts whatever unemployment emerges. Under Belief 3, the government thinks that there is a short-term trade-off between inflation and unemployment when $\gamma_1 < 0$, but also that there is no 'long-term' trade-off. Here the government believes that it can manipulate the positions of future Phillips curves (see Sargent (1999)). When δ is close to 1, the Phelps problem eventually tells the government to set inflation close to zero.

The preceding discussion is incomplete because it proceeds as though the government's beliefs were free parameters. However, in a self-confirming equilibrium they are outcomes, not inputs into the analysis.

2.4. Self-confirming equilibrium. Define $\xi_n = [W_{1n} \ W_{2n} \ X'_{n-1}]'$ and

$$(2.9) \quad g(\gamma, \xi_n) = \eta_n \begin{bmatrix} \pi_n \\ X_{n-1} \end{bmatrix}.$$

Notice that $g(\gamma, \xi_n)$ is the time n value of the object whose expectation is set to zero by the following orthogonality conditions for choosing γ to maximize the fit of (2.6):

$$(2.10) \quad \bar{g}(\gamma) \equiv E g(\gamma, \xi_n) = 0.$$

Equations (2.10) are the orthogonality conditions that identify γ in (2.6) as the population least-squares regression coefficients. Condition (2.12) thus renders the government's beliefs consistent with the data.

Let W^n denote the history of the joint shock process $\begin{bmatrix} W_{1n} \\ W_{2n} \end{bmatrix}$ up to n . Evidently, from (2.2), (2.3), (2.4), (2.8), X_{n-1} and therefore the ξ_n process are both functions of γ :

$$(2.11) \quad \xi_n = \xi(\gamma, W^n).$$

Via the least squares orthogonality conditions (2.10), γ depends on the stochastic process for ξ_n . But through the government's decision rule $h(\gamma)X_{n-1}$, ξ_n depends on γ . This sets up a mapping from γ to γ , a fixed point of which (denoted $\bar{\gamma}$) is a self-confirming equilibrium.

Definition 2.1. A self-confirming equilibrium (SCE) is a $\bar{\gamma}$ and a probability distribution for $(U_n, \pi_n, x_n, \hat{x}_n)$ that satisfies (2.2), (2.3), (2.4), (2.8) and for which

$$(2.12) \quad \bar{g}(\bar{\gamma}) = Eg(\bar{\gamma}, \xi_n) = 0.$$

The expectation in (2.12) is taken with respect to the probability distribution generated for (U_n, π_n, x_n) .

2.5. The self-confirming equilibrium outcome equals the time-consistent outcome. A simple calculation shows that Belief 1 supports a self-confirming equilibrium and that the self-confirming equilibrium outcome is $x_n = u$. This outcome is identical with the time-consistent outcome of Kydland and Prescott (1977), which emerges when the government knows the true model but confronts a timing protocol that prevents it from committing to a superior zero-inflation policy $x_n = 0$. If the government knows the true model and can commit, the zero inflation policy $x_n = 0$ emerges as the *Ramsey outcome*. Kydland and Prescott (1977) asserted that the timing protocol leading to the time-consistent outcome was the natural one. The inferiority of time-consistent outcome vis a vis the Ramsey outcome has come to be known as the time-inconsistency problem. The time inconsistency problem stems from the government being laden not with an incorrect model or with inappropriate objectives, but rather having a limiting timing protocol.

Notice that if the government has an incorrect model of the type we have assumed, then Beliefs 2 and 3 do support either the Ramsey outcome or a good approximation of it. However, neither of these beliefs is consistent with a self-confirming equilibrium.

2.6. Adaptation. The spirit of modern adaptive models is modestly (and perhaps only temporarily) to retreat from a self-confirming equilibrium by replacing (2.8) with $x_n = h(\gamma_n)X_{n-1}$, where $h(\cdot)$ is the *same* government best response function, and γ_n is a real-time estimator of the population regression coefficients γ , based on the history of observations on U_n, π_n . A government that substitutes its least-squares estimates γ_n for γ in (2.8) is said to be *adaptive*. A folk-theorem asserts that if a system in which adaptive decision makers employ a least squares learning scheme converges, then it converges to a self-confirming equilibrium. The folk-theorem holds true in our setting, but there is more to say. In particular, simulations by Sargent (1999) and Cho, Williams, and Sargent (2002) show that while a strong force propels an adaptive system toward a self-confirming equilibrium, there also occur recurrent and dramatic *escapes* that interrupt convergence to the self-confirming Nash equilibrium inflation outcome and that for a while make the adaptive government put inflation at or near its time-inconsistent optimal (Ramsey) value. After those escapes, the government's beliefs approximate either Belief 2 or 3, prompting the government to respond with Ramsey settings for so long as those beliefs persist. In the next section, we prepare to analyze how drift in the government's beliefs affects outcomes.

3. DRIFTING GOVERNMENT BELIEFS

We modify the estimator in Cho, Williams, and Sargent (2002) but leave the rest of their model intact. We motivate the government's estimator explicitly in terms of a Bayesian

prior that describes how coefficients in its econometric model drift. In particular, we attribute to the government the view that the data are generated by:

$$(3.13) \quad U_n = \alpha'_{n-1} \Phi_n + \eta_n$$

$$(3.14) \quad \alpha_n = \alpha_{n-1} + \Lambda_n$$

where $E\eta_n^2 = \sigma^2$, $\text{cov}(\Lambda_n) = V \ll \sigma^2 I$, and $\Phi_n \equiv \begin{pmatrix} \pi_n \\ X_{n-1} \end{pmatrix}$. The government's estimator is $\gamma_n \equiv \hat{\alpha}_{n|n-1}$, its optimal estimate of α_n conditional on information up to date $n-1$. It is well known that the optimal estimates in this linear model are provided by the Kalman filter. The Kalman filtering equations are:

$$(3.15) \quad \gamma_{n+1} = \gamma_n + \frac{P_n}{1 + \Phi_n^T P_n \Phi_n} \Phi_n (U_n - \Phi_n^T \gamma_n)$$

$$(3.16) \quad P_{n+1} = P_n - \frac{P_n \Phi_n \Phi_n^T P_n}{1 + \Phi_n^T P_n \Phi_n} + \sigma^{-2} V.$$

Here $\text{cov}(\gamma_n - \alpha_n) \equiv \sigma^2 P_n$. Let $M = M(\gamma) = E\Phi\Phi^T$, where we allow this moment matrix to depend on γ because the government's beliefs affect its decisions and therefore the moments M . Benveniste, Métivier, and Priouret (1990) note that for large n (3.16) is well approximated by:

$$(3.17) \quad P_{n+1} = P_n - P_n M(\gamma_n) P_n + \sigma^{-2} V.$$

Using the approximation in (3.17), we propose to replace the government's estimator (3.21)-(3.22) in CWS with:

$$(3.18) \quad \gamma_{n+1} = \gamma_n + P_n \Phi_n (U_n - \Phi_n^T \gamma_n)$$

$$(3.19) \quad P_{n+1} = P_n - P_n M(\gamma_n) P_n + \sigma^{-2} V.$$

In simulation results for our application, the exact Kalman filter (3.15)-(3.16) and the approximation (3.18)-(3.19) led to nearly indistinguishable results. Thus for the remainder of the paper we focus on the approximation, whose analysis proved more direct. Note that for fixed γ , P_n converges to a constant matrix P that solves an algebraic Riccati equation:

$$(3.20) \quad -PM(\gamma)P + \sigma^{-2}V = 0.$$

To get the constant gain algorithm, CWS implicitly assumed that $V = \epsilon^2 \sigma^2 M(\gamma)^{-1}$ and further that the standard deviation σ of the regression error in (3.13) is equal to the true standard deviation σ_1 of the unemployment shocks in (2.2). In this case, it is easy to see from (3.20) that the limiting value is $P = \epsilon M(\gamma)^{-1}$. The algorithm corresponding to (3.15), (3.16) becomes

$$(3.21) \quad \gamma_{n+1} = \gamma_n + \epsilon R_n^{-1} \Phi_n (U_n - \Phi_n^T \gamma_n)$$

$$(3.22) \quad R_{n+1} = R_n + \epsilon (\Phi_n \Phi_n^T - R_n),$$

the constant gain algorithm assumed by CWS. This is often referred to as recursive least squares (RLS). Note that this has the same asymptotic behavior of (3.15) and (3.16) under the assumed form of V , since R_n converges to $M(\gamma)$. However the transient responses may differ, as we show below.

As CWS did for a constant gain algorithm, here we want to study what happens as the rate of variation in the parameters declines to zero. In particular, we will suppose that the covariance matrix of the parameters takes the form $V = \epsilon^2 \hat{V}$. We will study limits as $\epsilon \rightarrow 0$, and to denote the limits we use an ϵ superscript on the parameters. From (3.17) and (3.20) we see that $P_n^\epsilon \rightarrow 0$ as $\epsilon \rightarrow 0$. Therefore we define a scaled matrix $\hat{P}_n^\epsilon = \epsilon P_n^\epsilon$ that does not vanish. Using these scalings, we can re-write (3.18)-(3.19) as:

$$(3.23) \quad \gamma_{n+1}^\epsilon = \gamma_n^\epsilon + \epsilon \hat{P}_n^\epsilon \Phi_n (U_n - \Phi_n^T \gamma_n)$$

$$(3.24) \quad \hat{P}_{n+1}^\epsilon = \hat{P}_n^\epsilon + \epsilon \left(\sigma^{-2} \hat{V} - \hat{P}_n^\epsilon M(\gamma_n^\epsilon) \hat{P}_n^\epsilon \right).$$

Then under conditions analogous to those specified in Cho, Williams, and Sargent (2002), as $\epsilon \rightarrow 0$ the sequence $\{\gamma_n^\epsilon, \hat{P}_n^\epsilon\}$ converges weakly to the solution of the following ODEs:

$$(3.25) \quad \dot{\gamma} = P\bar{g}(\gamma)$$

$$(3.26) \quad \dot{P} = \sigma^{-2} \hat{V} - PM(\gamma)P.$$

If the ODEs have a stable point $(\bar{\gamma}, \bar{P})$, we can expect that beliefs will converge to it as $\epsilon \rightarrow 0$. The limiting beliefs $\bar{\gamma}$ comprise a self-confirming equilibrium.

4. RELATION OF KALMAN FILTER TO RLS

In this section we compare the behavior of our generalized learning rule to the one used by CWS. We have argued that, under appropriate conditions, as $\epsilon \rightarrow 0$ the sequence $\{\gamma_n^\epsilon, \hat{P}_n^\epsilon\}$ converges weakly to the solution of the ODEs (3.25)-(3.26). Under the RLS learning rule (3.21)-(3.22) used by CWS, the corresponding ODEs are:

$$(4.27) \quad \dot{\gamma} = R^{-1} \bar{g}(\gamma)$$

$$(4.28) \quad \dot{R} = M(\gamma) - R.$$

4.1. RLS and a special case of the Kalman filter. Constant gain recursive least squares learning is often justified by noting its relationship with the Kalman filter in a setting where beliefs drift. However, the relationship is only asymptotic – for large samples the two algorithms will give the same result, but the small sample properties may differ. A way of highlighting this distinction is to note that the RLS differential equations (4.27)-(4.28) and the corresponding special case of the Kalman filter ODEs (3.25)-(3.26) have the same asymptotic behavior but slightly different transient behavior. To illustrate this, we initialize the ODEs at $\gamma = [5, 0]'$, and we set P and R at their equilibrium points. Figure 1 shows the time paths of the regression coefficients (the γ), and Figure 2 shows elements of the P vector from the Kalman filter and the corresponding R^{-1} vector from the constant gain RLS. The coefficients converge to the same stable points, but the Kalman filter converges much faster.

These differences are also evident in the discrete time simulations of the learning rules. In Figure 3 we plot the time paths of inflation that result from the different learning rules. The top panel plots the CG specification (3.21)-(3.22), while the bottom panel plots the Kalman filter specification (3.15)-(3.16). (Nearly identical results occurred when we used approximation (3.18)-(3.19).) The self-confirming equilibrium inflation level in this specification is 5. Later sections of the paper will describe the episodes evident in




FIGURE 1. Solution of mean dynamics ODE for the regression coefficients, under Kalman filter and constant gain RLS algorithms.




FIGURE 2. Solution of mean dynamics ODE for elements of the P matrix, under Kalman filter and constant gain RLS algorithms.

the figure in which the inflation rate falls dramatically from the SCE level, because the government’s beliefs “escape” from the SCE. Here we focus on the differences in the convergence to the SCE following the escape to low inflation levels. As the figure shows, inflation increases much more rapidly under Kalman filter learning than under RLS. This




FIGURE 3. Simulated time series of inflation under the constant gain RLS and Kalman filter algorithms.

agrees with our results above. We defer until later a discussion of the escape dynamics, but for now note that the escape dynamics are very similar in the two cases.

4.2. Stability. Thus far we have focused on the differences between RLS and the corresponding special case of the Kalman filter. We now broaden the scope of our inquiry to consider the Kalman filter under an arbitrary specification of the parameter drift covariance matrix V . An interesting difference between Kalman filter learning and RLS learning in this case is the possibility of cycles under the Kalman filter. We explore this in more detail in the next section, but here provide the conditions insuring stability of a steady state. The conditions for stability here extend the “E-Stability” conditions discussed in Evans and Honkapohja (2001), where now prior beliefs about the form of parameter drift matter for the stability of the SCE. Similar extensions in a different context are provided by Honkapohja and Mitra (2002).

In the RLS case, we can determine the stability of a steady state $\bar{\gamma}$ solely by analyzing the Jacobian $\frac{\partial \bar{q}}{\partial \gamma}(\bar{\gamma})$. Thus stability in this case is governed by the E-stability conditions as in Evans and Honkapohja (2001). First we introduce some more notation. Define $\theta = [\gamma', \text{col}(R)']'$ and the stacked right side of (4.27)-(4.28) as $\bar{b}^{RLS}(\theta)$. Evaluated at $\bar{\theta} = (\bar{\gamma}, \bar{R})$, the Jacobian of $\bar{b}^{RLS}(\theta)$ has the block form:

$$(4.29) \quad \bar{b}_{\theta}^{RLS}(\bar{\theta}) = \begin{bmatrix} \bar{R}^{-1} \frac{\partial \bar{q}}{\partial \gamma}(\bar{\gamma}) & 0 \\ \text{col}(\frac{\partial M}{\partial \gamma}(\bar{\gamma})) & -I \end{bmatrix}$$

There are four eigenvalues of -1 , and the eigenvalues of $\frac{\partial \bar{q}}{\partial \gamma}(\bar{\gamma})$ determine the stability. The repeated eigenvalues of the lower block imply that the eigenspace of the lower block is a subset of the eigenspace of the upper block. This implies that stability is determined by

$\frac{\partial \bar{g}}{\partial \gamma}(\bar{\gamma})$, independent of \bar{R} . For completeness, note that in the static version of our model, in which γ contains only an intercept and a slope, this Jacobian is:

$$(4.30) \quad \frac{\partial \bar{g}}{\partial \gamma}(\bar{\gamma}) = - \begin{bmatrix} 0.5 & u \\ 0.5u & u^2 + \sigma_2^2 \end{bmatrix}.$$

It is then evident that both eigenvalues of (4.30) are real and negative, and thus the SCE $\bar{\gamma}$ is always (at least locally) stable.

However under the Kalman filter, we cannot make this reduction. Now let $\theta = [\gamma', \text{col}(\bar{P})]'$ and define the stacked right side of (3.25)-(3.26) as $\bar{b}(\theta)$. The Jacobian in this case takes the form:

$$(4.31) \quad \bar{b}_\theta(\bar{\theta}) = \begin{bmatrix} \bar{P} \frac{\partial \bar{g}}{\partial \gamma}(\bar{\gamma}) & 0 \\ \text{col}(-\bar{P} \frac{\partial \bar{M}}{\partial \gamma}(\bar{\gamma}) \bar{P}) & -\text{col}(\frac{\partial}{\partial P} \bar{P} M \bar{P}) \end{bmatrix}.$$

Since PMP is quadratic in P , the lower right block always has eigenvalues with negative real part. However in general, unlike the RLS special case, the eigenspace of this lower block is a *superset* eigenspace of the upper block. Therefore to determine the stability of a steady state, we must examine the eigenvalues of $\bar{P} \frac{\partial \bar{g}}{\partial \gamma}(\bar{\gamma})$, not just $\frac{\partial \bar{g}}{\partial \gamma}(\bar{\gamma})$ as before.⁶

Further, some specifications of V make the SCE $\bar{\gamma}$ unstable. This happens by flipping the sign of $\text{trace}(\bar{P} \frac{\partial \bar{g}}{\partial \gamma}(\bar{\gamma}))$ from negative to positive. Using (4.30) and the symmetry of $\bar{P} = \begin{bmatrix} P_{11} & P_{12} \\ P_{12} & P_{22} \end{bmatrix}$, the steady state becomes unstable when:

$$-1.5uP_{12} > 0.5P_{11} + (u^2 + \sigma_2^2)P_{22}.$$

Since \hat{V} determines \bar{P} through an analogue of the Riccati equation (3.20), this reduces to a restriction on \hat{V} .

5. BIFURCATION AND CYCLES

In this section we show that the change of the stability of the steady state produces a unique stable limit cycle. This is an example of a Hopf bifurcation, yielding what Bullard (1994) calls “learning equilibria.” That is, even though the underlying model does not have cycles, the introductions of learning produces them. Although the full vector differential equation (3.25)-(3.26) cycles, it suffices to analyze the differential equation governing the regression coefficients γ . Therefore we avail ourselves of well-developed methods for analysis of ODEs in the plane. In particular, we have the following result.

Theorem 5.1. *Let \bar{P} be the stable solution of (3.26) evaluated at $\bar{\gamma}$. Then if $\text{trace}(\bar{P} \frac{\partial \bar{g}}{\partial \gamma}(\bar{\gamma})) < 0$, the unique steady state $\bar{\gamma}$ is locally stable. A unique stable limit cycle bifurcates from $\bar{\gamma}$ as $\text{trace}(\bar{P} \frac{\partial \bar{g}}{\partial \gamma}(\bar{\gamma}))$ increases from zero.*

The theorem follows from Theorem 1 and Remark 1, pp. 343-344 of Perko (1996). In particular, we must reparametrize the system by centering around the steady state $\bar{\gamma}$.

⁶Since \bar{P} is symmetric and positive definite, in the terminology of Arrow and McManus (1958) we thus require that $\frac{\partial \bar{g}}{\partial \gamma}(\bar{\gamma})$ be S -stable. See also Honkapohja and Mitra (2002).

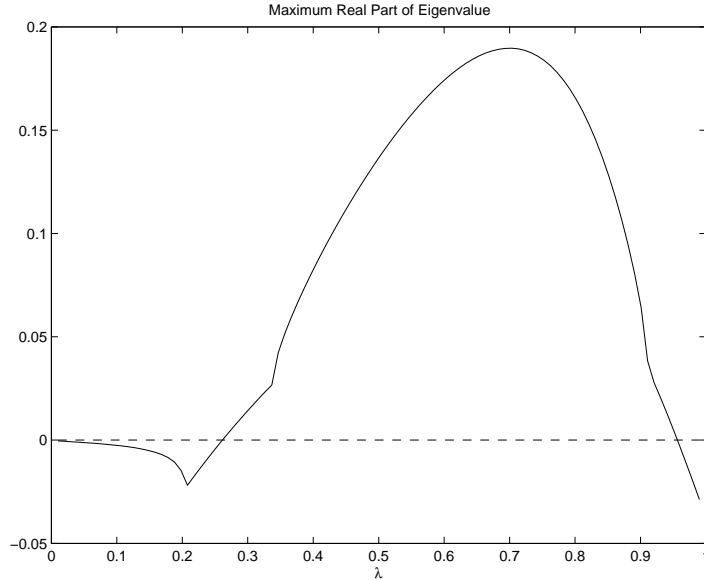


FIGURE 4. The maximum real part of the eigenvalues of $\bar{P} \frac{\partial \bar{g}}{\partial \gamma}(\bar{\gamma})$ where $V = V(\lambda)$ for different λ .

Then the conditions of the theorem are stated in terms of the linearized evolution. A limit cycle emerges when the real parts of the eigenvalues change signs at the steady state.⁷

We actually discovered this bifurcation by exploring simulations with different values of the prior V . In particular, we examined cases where we started at the prior associated with RLS learning, $V^* = \epsilon^2 \sigma^2 M$. We tightened the priors on the slope coefficient by setting $V = V(\lambda)$ for different λ , where:

$$V(\lambda) = \begin{bmatrix} V_{11}^* & \sqrt{\lambda} V_{12}^* \\ \sqrt{\lambda} V_{12}^* & \lambda V_{22}^* \end{bmatrix}.$$

As λ decreases from one, over an intermediate range the stability of the steady state flips, as $\text{trace}(\bar{P} \frac{\partial \bar{g}}{\partial \gamma}(\bar{\gamma}))$ changes signs. This is shown in Figure 4, which plots the maximum real part of the eigenvalues of $\bar{P} \frac{\partial \bar{g}}{\partial \gamma}(\bar{\gamma})$. As the theorem suggests, this leads to a unique stable limit cycle. This is shown in Figure 5, which plots the time paths of the coefficients from (3.25)-(3.26), along with their phase plot for the case of $\lambda = 0.7$. Here we initialize the ODEs near $\bar{\gamma}$ (which is shown with an “x” in the right plot). The figure clearly shows the cycle that emerges in the mean dynamics, in which beliefs start near the SCE but move away from it and continually oscillate.

Thus for some specifications of prior beliefs, we find that the mean dynamics lead to cyclical fluctuations. Along these cycles, the inflation rate oscillates from near the Nash level to near the Ramsey level, as Figure 6 shows. The top panel of the figure shows a simulated time series of inflation, while the bottom panel plots for comparison the mean dynamics of the slope coefficient. Instead of being driven by rare sequences of stochastic

⁷There is a further qualification on the sign of the Lyapunov number, which depends on the second and third order terms in the Taylor expansion of $\bar{g}(\gamma)$, which holds here.




FIGURE 5. A stable limit cycle that bifurcates from the stable steady state.

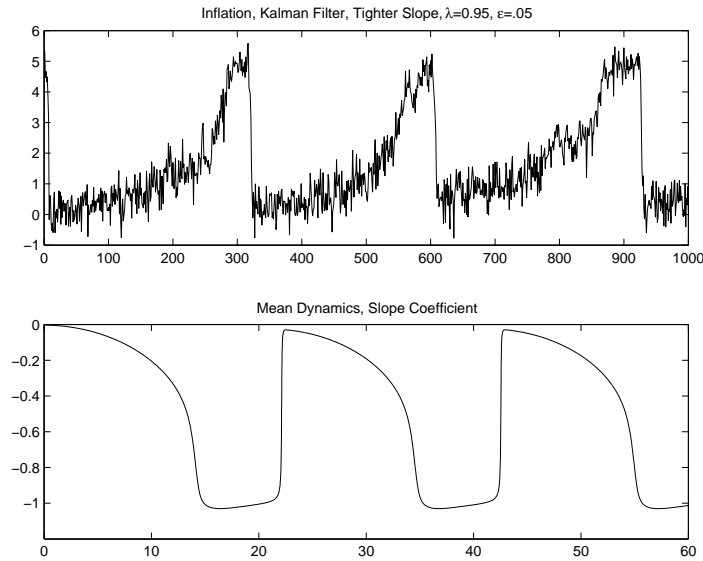


FIGURE 6. Simulated time series of inflation and the mean dynamics of the slope coefficient along a limit cycle.

shocks which cause agents to escapes from the SCE, here the disinflationary episodes come from the mean dynamics themselves and so are a typical feature of the time series. While in the rest of the paper we study situations in which the SCE is stable, the cycles which we have discovered here already suggest that the mean dynamics have some interesting

features that directly bear on the escape dynamics. Loosely speaking, even in the stable case the mean dynamics exhibit some of the “roundabout” dynamics that are evident in the cycles.⁸ The key difference is that as $\epsilon \rightarrow 0$ the disinflations persist in a limit cycle, while they become increasingly rare in the stable case.

6. EFFECT OF PRIORS ON THE SPEED OF CONVERGENCE

We now continue our analysis of the effects of changes in the prior covariance matrix for the parameter drift V . We have already seen that some specifications of this matrix can lead to limit cycles, but now we focus on cases in which there are no limit cycles. In these cases, the prior covariance matrix will affect the dynamics of beliefs, but not their asymptotic behavior. However different V matrices may imply substantially different time series behavior.

What matters most is how changes in the V matrix alter the mean dynamics. However, it is also useful to apply results on the *speed* of convergence to a steady state. These results extend the *ordinary* differential equation characterization in (3.25)-(3.26) to a *stochastic* differential equation (SDE). Again a bit more notation. Let $\xi_n = [W'_n, \Phi'_n]'$ and define the stacked version of (3.23)-(3.24) as:

$$\theta_{n+1} = \theta_n + \epsilon b(\theta_n, \xi_n).$$

Note that above we have defined $\bar{b}(\theta) = E[b(\theta, \xi_n)]$. Then under appropriate conditions, Benveniste, Métivier, and Priouret (1990) and Kushner and Yin (1997) show that as $\epsilon \rightarrow 0$ we can approximate the time paths of (3.18)-(3.19) by the SDE:

$$(6.32) \quad d\theta^\epsilon(t) = b'_\theta(\bar{\theta}) (\theta^\epsilon(t) - \bar{\theta}) dt + \sqrt{\epsilon} \Sigma_\theta dB(t),$$

where $\Sigma_\theta \Sigma'_\theta = \text{cov}[b(\theta, \xi_n)]$.⁹

Notice that the evolution of the P block in (3.19) is a deterministic function of γ and P . This implies that Σ_θ is nonzero only in the upper block corresponding to γ . Also note that the covariance matrix of the γ coefficients takes the following form:

$$(6.33) \quad Q(\gamma, P) = P \hat{Q}(\gamma) P \equiv P \text{cov}[(U_n - \Phi'_n \gamma) \Phi_n] P,$$

where it helps to recall that both U and Φ depend on γ . Note that by (3.20), in a self-confirming equilibrium we have the following equivalences:

$$\hat{Q}(\bar{\gamma}) = \sigma^2 M(\bar{\gamma}), \text{ so } Q(\bar{\gamma}, \bar{P}) = \hat{V}.$$

In other words, at the SCE all variation in the parameters is due to the prior beliefs. Decomposing this variation by defining $\Sigma \Sigma' = \hat{V}$, we see that (6.32) implies the following SDE for the parameters:

$$(6.34) \quad d\gamma^\epsilon(t) = \bar{P} \frac{\partial \bar{g}}{\partial \gamma}(\bar{\gamma}) (\gamma^\epsilon(t) - \bar{\gamma}) dt + \sqrt{\epsilon} \Sigma dB(t).$$

⁸In the stable case, the SCE is locally stable with a certain radius of stability. Outside this radius, the mean dynamics point away from the SCE toward the Ramsey outcome, although they eventually lead back to the SCE (which is thus asymptotically stable). These features were noticed by Evans and Honkapohja (2001) and analyzed by Cho, Williams, and Sargent (2002). See also Section 7.2 below.

⁹See Williams (2001) for a verification of the necessary conditions in a closely related model.

For scalar diffusions, a measure of the speed of pull back to the mean is given by the instantaneous mean (the drift) divided by the instantaneous standard deviation. We use a similar approach here to provide a local measure of the strength of pull, which we define as:

$$\begin{aligned}\rho(\gamma) &= \Sigma^{-1} \bar{P} \frac{\partial \bar{g}}{\partial \gamma}(\bar{\gamma})(\gamma - \bar{\gamma}) \\ &= \hat{Q}(\bar{\gamma})^{-1/2} \frac{\partial \bar{g}}{\partial \gamma}(\bar{\gamma})(\gamma - \bar{\gamma}),\end{aligned}$$

where the second equality uses (6.33). Notice that since $\hat{Q}(\bar{\gamma})$ is independent of \hat{V} , the local speed of convergence is independent of the prior specification V . The priors affect the local evolution of beliefs in (6.34) through both the mean (through \bar{P}) and the standard deviation (Σ is the matrix square root of \hat{V}). In both cases the effects are linear, so their ratio is unaffected.

Prior beliefs about parameter drift thus do not affect the local speed of convergence, although they may have higher order or non-local effects. For example, the differences between the Kalman filter and RLS that we observed in the simulations are non-local (as we initialized the beliefs far away from the SCE). Note that a counterpart to (6.32) holds for the RLS case, with a slight correction due to variation in the R matrix. But since \bar{g} and \hat{Q} are the same in both cases, our measure of pull ρ will be nearly the same in both cases. Thus the ρ measure does not characterize the relative speeds. Instead, the difference between the algorithms is driven by how the P matrix in the Kalman case and the R matrix in the RLS case respond to changes in beliefs γ . This is governed by $PM_\gamma(\gamma)P$ for the Kalman case and $M_\gamma(\gamma)$ for the RLS case, which are elements of \bar{b}_θ and \bar{b}_θ^{RLS} that mainly drive the drift of P and R .

7. EFFECT OF PRIORS ON ESCAPE DYNAMICS

7.1. General Results. We now turn toward analyzing escapes, which we define as events in which the government's beliefs move some fixed distance away from the SCE. Cho, Williams, and Sargent (2002) apply the results of Williams (2002) to characterize the escape dynamics. We apply those same methods here. With the present specification of beliefs, the escape dynamics from a set G can be characterized by the following control problem on some finite horizon \bar{T} :

$$(7.35) \quad \bar{S} = \inf_{v(\cdot), T} \frac{1}{2} \int_0^T v(s)' Q(\gamma(s), P(s))^\dagger v(s) ds$$

where Q^\dagger is the pseudoinverse of Q , and the optimization is subject to:

$$(7.36) \quad \dot{\gamma} = P\bar{g}(\gamma) + v$$

$$(7.37) \quad \dot{P} = \sigma^{-2} \hat{V} - PM(\gamma)P$$

$$(7.38) \quad \gamma(0) = \bar{\gamma}, P(0) = \bar{P}, \gamma(T) \in \partial G \text{ for some } 0 < T \leq \bar{T}.$$

We will focus on spherical sets G with radius r :

$$G(r) = \{\gamma : (\gamma - \bar{\gamma})'(\gamma - \bar{\gamma}) \leq r^2\}$$

In general dynamic models, the key matrix Q is determined by solving some Lyapunov equations. However in the static version of our model, Q was defined in (6.33) above and is explicitly calculable. For any specification of \hat{V} , we can solve this variational problem to determine the rate function \bar{S} and the dominant escape path $\gamma(\cdot)$. In general, numerical methods are needed to solve the problem; we show some numerical results in Section 7.4 below. However, we can get some additional insight on various aspects of the problem in the following sections that provide some approximation results.

The following theorem is established (under some regularity conditions) in Williams (2002). It states precisely the sense in which the control problem characterizes the escape dynamics.

Theorem 7.1. *Fix a horizon $\bar{T} < \infty$ and a set $G \subset \mathcal{G}$ with $\bar{\gamma} \in G$. Let $\gamma^\varepsilon(\cdot)$ be the piecewise linear interpolation of the belief sequence $\{\gamma_n^\varepsilon\}$, and let $\gamma(\cdot) : [0, \bar{T}] \rightarrow \mathbb{R}^p$ solve (7.35), attaining the value \bar{S} , with $\gamma(\bar{T}) \in \partial G$ for $T \leq \bar{T}$.*

1. *Suppose that the shocks W_n are i.i.d. and unbounded. Then we have:*

$$\limsup_{\varepsilon \rightarrow 0} \varepsilon \log P(\gamma^\varepsilon(t) \notin G \text{ for some } 0 < t \leq \bar{T} | \gamma^\varepsilon(0) = \bar{\gamma}) \leq -\bar{S}.$$

2. *Suppose that the shocks W_n are i.i.d. and bounded, and \bar{S} is continuous as a function of the radius of the set G . Then we have:*

$$\lim_{\varepsilon \rightarrow 0} \varepsilon \log P(\gamma^\varepsilon(t) \notin G \text{ for some } 0 < t \leq \bar{T} | \gamma^\varepsilon(0) = \bar{\gamma}) = -\bar{S}.$$

3. *Under the assumptions of part 2, for all $\delta > 0$:*

$$\lim_{\varepsilon \rightarrow 0} P[\exp((\bar{S} + \delta)/\varepsilon) > \tau^\varepsilon > \exp((\bar{S} - \delta)/\varepsilon)] = 1,$$

$$\text{and } \lim_{\varepsilon \rightarrow 0} \varepsilon \log E(\tau^\varepsilon) = \bar{S}.$$

4. *Under the assumptions of part 2, for any $\gamma^\varepsilon(\tau^\varepsilon)$ and $\delta > 0$:*

$$\lim_{\varepsilon \rightarrow 0} P(|\gamma^\varepsilon(\tau^\varepsilon) - \gamma(\bar{T})| < \delta) = 1.$$

7.2. Effects on Direction of Escape. While the full dynamic optimization problem (7.35) is rather computationally intensive, we can find the direction of escape in some special cases. Fully to characterize the escape dynamics, we are interested in the complete time paths of beliefs as well as the value of the minimized cost function \bar{S} . However some insight is provided by examining the directions in which beliefs move during an escape. Williams (2002) shows that the instantaneous escape directions can be computed by solving a static minimization problem.

Taking limits as \bar{T} goes to zero for the escape set $G(r)$ gives the static problem:

$$(7.39) \quad \inf_v \frac{1}{2} v' Q(\bar{\gamma}, \bar{P})^\dagger v \text{ subject to } \|v\| = r.$$

The first order condition for (7.39), with Lagrange multiplier $\mu/2$ on the constraint, is:

$$(7.40) \quad Q(\bar{\gamma}, \bar{P})^\dagger v^* = \mu v^*, \text{ or: } Q(\bar{\gamma}, \bar{P}) v^* = \frac{1}{\mu} v^*.$$

Thus the minimizing static escape direction v^* is the eigenvector associated with the largest eigenvalue of $Q(\bar{\gamma}, \bar{P})$ (and so the smallest of Q^\dagger), normalized so that $\|v^*\| = r$.

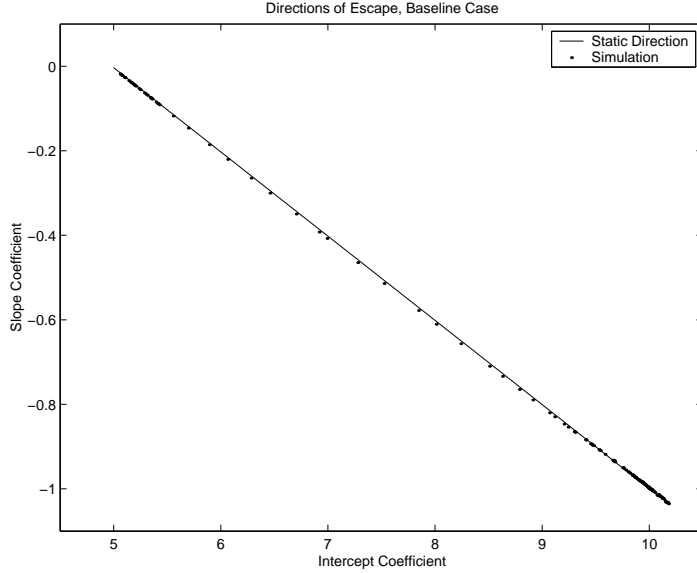


FIGURE 7. The escape direction in the baseline case.

Notice that the escape may occur in either the positive or negative direction along this eigenvector, so that the terminal points of the dominant paths are $\bar{\gamma} \pm v^* \in \partial G(r)$. The instantaneous escapes use no information about the evolution of beliefs, and so inherit the symmetry of the cost function. We show below that nonlinearities in the belief evolution make escapes in one direction more likely than the other. We also show how to choose the radius r , which has thus far been arbitrary, in a meaningful way.

We noted above that $Q(\bar{\gamma}, \bar{P}) = \hat{V}$, the limiting covariance matrix of the priors. This suggests that the instantaneous escape directions for different specifications \hat{V} can be found by examining the eigenvectors. While the eigenvectors can be found for any given value of the noise σ , we obtain our most explicit analytic results in the limit as $\sigma \rightarrow 0$. The following results directly parallel those in Williams (2002) and consider the baseline case where analogous to RLS, with $\hat{V} = V^* = \sigma^2 M$. We also focus on the static belief specification, in which the government's regression contains no lags. Under this specification, letting $\hat{r} = r/(u^2 + 1)$, we see that v^* from (7.40) satisfies:

$$(7.41) \quad \lim_{\sigma \rightarrow 0} v^* = \hat{r} \begin{bmatrix} -u \\ 1 \end{bmatrix}.$$

The terminal points of the instantaneous escape paths satisfy:

$$\lim_{\sigma \rightarrow 0} (\bar{\gamma} + v^*) = \begin{bmatrix} u(2 - \hat{r}) \\ -1 + \hat{r} \end{bmatrix}.$$

For $\hat{r} = 1$, the terminal point is $[u, 0]'$, which is the Belief 2 defined in Section 2 above, and which supports the Ramsey outcome. Therefore, for small σ we expect instantaneous escapes to occur in the direction of the Ramsey outcome.

In general, escapes from any but the smallest sets will not happen instantaneously. But as $\sigma \rightarrow 0$ large escapes do happen arbitrarily fast. For any $\sigma > 0$, the SCE $\bar{\gamma}$ is a locally stable limit point of the mean dynamics (3.25)-(3.26). Thus after a small perturbation, agents return to the SCE. Formally, there exists some $\delta > 0$ such that:

$$\langle \bar{P}g(\bar{\gamma} + v), v \rangle < 0, \forall v \in \mathbb{R}^2 \text{ s.t. } \|v\| \leq \delta.$$

Therefore we require “large enough” nonzero perturbations v in order to escape from the SCE. But as $\sigma \rightarrow 0$ the radius of local stability $\delta \rightarrow 0$. In fact, we have:

$$\lim_{\sigma \rightarrow 0} \langle \bar{P}g(\bar{\gamma} + v^*), v^* \rangle = u \frac{\hat{r}(1 - \hat{r})}{1 + (1 - \hat{r}^2)} (x(\bar{\gamma} + v^*) - u) (u^2 + 1) = \begin{cases} 0 & \text{for } \hat{r} = 0 \text{ or } \hat{r} = 1 \\ > 0 & \text{for } 0 < \hat{r} < 1 \end{cases}$$

Further, although the instantaneous escapes can occur in either positive or negative directions, along the negative direction we have:

$$\lim_{\sigma \rightarrow 0} \langle \bar{P}g(\bar{\gamma} - v^*), -v^* \rangle = -u \frac{\hat{r}(1 + \hat{r})}{1 + (1 - \hat{r}^2)} (x(\bar{\gamma} - v^*) - u) (u^2 + 1) < 0 \text{ for } \hat{r} > 0.$$

For perturbations along v^* in a positive direction, the mean dynamics point away from the SCE, while perturbations in the negative direction lead back to the SCE.

As $\sigma \rightarrow 0$, the SCE becomes unstable and so the necessary escapes become arbitrarily small and happen arbitrarily fast. Once beliefs start along v^* in a positive direction, they can continue in this way “for free” with $v = 0$. Further, beliefs move along the same constant direction since:

$$\lim_{\sigma \rightarrow 0} \bar{P}g(\bar{\gamma} + v^*) = u \frac{\hat{r}(1 - \hat{r})}{1 + (1 - \hat{r}^2)} (x(\bar{\gamma} + v^*) - u) \begin{bmatrix} -u \\ 1 \end{bmatrix},$$

which clearly lies along the instantaneous direction. This continues all the way until $\hat{r} = 1$, which attains Belief 2 supporting the Ramsey outcome. Similar results hold for small nonzero σ . In this case, the SCE is stable, but the radius of stability δ is small. These results are illustrated in Figure 7, which plots a simulated escape path from the model along with our static prediction v^* . Here we clearly see that the simulation lines up directly with our prediction.

When we move away from the baseline case, the analysis becomes a bit more complicated. For example, consider the case in which we decrease the prior correlation between the slope and intercept coefficients, but keep the norm of the matrix constant. Thus we set $V = V(\lambda)$ for different λ , where:

$$V(\lambda) = \lambda V^* + \sqrt{(1 - \lambda^2) \frac{\text{tr}(V^{*2})}{2}} I.$$

These transformations preserve the eigenvectors of V^* so that the same instantaneous escape directions characterize the escape dynamics. However the escape paths are no longer straight lines. The static direction determines the initial direction of escape, and the mean dynamics continue to point outward for small σ , but now the mean dynamics follow a curved arc away from the SCE. An example of this is shown in Figure 8, which plots simulated escape paths for $\lambda = 1$ (the baseline case as above) and $\lambda = 0.63$. The escape path for the lower correlation is clearly longer and more varied, and tends to follow




FIGURE 8. The escape direction with a lower correlation between the slope and intercept coefficients.




FIGURE 9. The escape direction with a tighter prior on the intercept coefficient.

a curved path to the boundary of the set. But the end point appears to be roughly the same as in the baseline case.

A different story emerges when we tighten the priors on the individual coefficients. We have already seen that tightening the prior on the slope coefficient can lead to limit cycles. Tightening the prior on the intercept coefficient does not lead to instability or cycles, but

it does alter the escape dynamics. In particular, we now set $V = V(\lambda)$ for different λ , where:

$$V(\lambda) = \begin{bmatrix} \lambda V_{11}^* & \sqrt{\lambda} V_{12}^* \\ \sqrt{\lambda} V_{12}^* & V_{22}^* \end{bmatrix}.$$

An illustration of this is shown in Figure 9 which plots simulated escape paths for $\lambda = 1$ (the baseline case as above) and $\lambda = 0.63$. For the tighter prior on the intercept, the escape path now starts in a curved arc away from $\bar{\gamma}$. It then follows a straight line with a different slope from the baseline case, leading to a different escape point.

Interestingly, the escapes all seem to lead to the zero inflation outcome, but they get there by different avenues and are supported by different beliefs.

7.3. Effects on Speed of Escape. Again, by solving the full dynamic optimization problem we can immediately see the effects of changes in the prior on the frequency of escapes (which is determined by \bar{S}). However to gain additional insight, it may be useful to consider some approximation results. In particular, here we apply the results of Mil'shtein and Ryashko (1995), who use perturbation arguments to characterize the rate function \bar{S} in a neighborhood of the stable point $\bar{\gamma}$. They present their results in the context of continuous time diffusions, but their methods apply directly to our setting. These results are also directly related to two of our approximation results above. Here we provide a quadratic approximation to the rate function, which would be exact if the belief evolution were linear. Thus our findings here are closely tied to the functional central limit theorem results above. Further, over arbitrarily short intervals, the quadratic approximation reduces to the static cost function that we used to calculate the instantaneous escape directions above.

Define $S(x)$ in an analogous fashion to \bar{S} in (7.35), but replace the boundary conditions (7.38) with:

$$(7.42) \quad \gamma(0) = \bar{\gamma}, \quad P(0) = \bar{P}, \quad \gamma(T) = x \text{ for some } 0 < T \leq \bar{T}.$$

Then $\bar{S} = \inf_{x \in \partial G} S(x)$. Mil'shtein and Ryashko (1995) show that the following approximation holds:

$$S(x) = \frac{1}{2}(x - \bar{\gamma})'W^{-1}(x - \bar{\gamma}) + O(|x - \bar{\gamma}|^3),$$

where W solves the Lyapunov equation:

$$\bar{P} \frac{\partial \bar{g}}{\partial \gamma}(\bar{\gamma})W + W \frac{\partial \bar{g}}{\partial \gamma}(\bar{\gamma})' \bar{P} + 2Q(\bar{\gamma}, \bar{P}) = 0.$$

Using the fact that $Q(\bar{\gamma}, \bar{P}) = \hat{V} = \sigma^2 \bar{P} M(\bar{\gamma}) \bar{P}$ we can simplify this to:

$$\frac{\partial \bar{g}}{\partial \gamma}(\bar{\gamma})Z + Z \frac{\partial \bar{g}}{\partial \gamma}(\bar{\gamma})' + 2\sigma^2 M(\bar{\gamma}) = 0,$$

where $Z = W \bar{P}^{-1}$. Therefore the terminal points of the escape paths are determined by the eigenvectors of W , just as in the instantaneous calculations above.

These expressions provide us with a convenient means of seeing how, to second order, changes in V change the speed of escape. The effect of the priors V comes through the change in \bar{P}^{-1} . While we saw above that our measure of pull ρ was unaffected by the priors, an alternative measure of pull used by Conley, Hansen, Luttmner, and Scheinkman (1997)

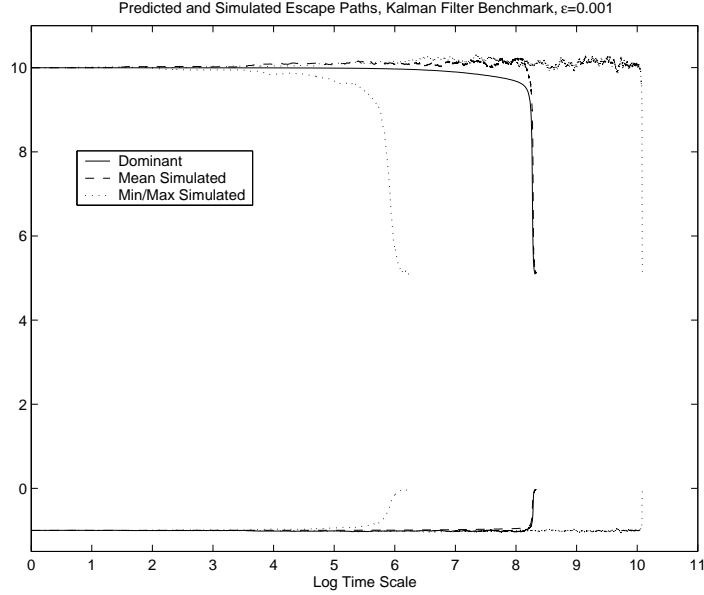


FIGURE 10. Predicted (dominant) escape path and the results from 1000 simulations.

is given the ratio of the instantaneous mean to the instantaneous variance (instead of the standard deviation). In our setting, this measure is linear in \bar{P}^{-1} .¹⁰ Settings of priors that make \bar{P}^{-1} “larger” lead to stronger pull back to the SCE. These settings also lead to larger values of $W^{-1} = Z^{-1}\bar{P}^{-1}$, which implies a greater exponential rate of decrease \bar{S} in the probability of escape. This relationship is intuitively clear: if beliefs are pulled back to the mean more strongly, it is more difficult to escape, and thus the likelihood of escapes decreases and the mean time between escapes increases.

7.4. Numerical Results. In this section we solve the full escape problem in (7.35) in order to determine the likelihood of escapes and obtain predictions about the most likely escape paths. Williams (2002) and Cho, Williams, and Sargent (2002) describe how to solve (7.35) numerically in the recursive least squares case. Modifications to let us handle the Kalman filter case are straightforward.

	Rate of Escape	Mean Escape Time				
		$\epsilon = .05$	$\epsilon = .01$	$\epsilon = .005$	$\epsilon = .002$	$\epsilon = .001$
Kalman	3.14×10^{-4}	22.9	166.6	412.1	2378.3	4138.1
CG RLS	4.99×10^{-4}	20.6	165.7	409.6	2414.8	4486.7

 TABLE 1. Rate of Escape (\bar{S}) and mean escape times for different gain settings ϵ in the Kalman filter benchmark and constant gain recursive least squares cases.

¹⁰This is not a measure of pull per unit time, but is instead related to the existence of a stationary distribution.

In the benchmark case where $\hat{V} = \sigma^2 M(\bar{\gamma})^{-1}$, which corresponds to the constant gain analysis of Cho, Williams, and Sargent (2002), we find very similar results as in that paper. We found that the escape directions lined up nearly exactly with those determined analytically above, and that our time paths are very similar to those Cho, Williams, and Sargent (2002) found in the constant gain case. We found that the rate of escape \bar{S} is slightly lower in the Kalman case than in the constant gain case. This is reflected in some simulations summarized in Table 1. There we see that the escape times for both cases increase exponentially as the gain decreases. While the mean escape times are very close, for very small gain settings the Kalman filter algorithm escapes faster than the recursive least squares algorithm. We also found that our predictions from our calculations provided a good job of matching up with the simulations, as illustrated in Figure 10. There we plot our calculated dominant escape path for the slope coefficient (bottom) and intercept coefficient (top), along with the results of 1000 simulated escapes for the very small gain setting of $\epsilon = 0.001$. The figure shows those paths with the minimum, maximum, and mean escape times. We clearly see that all paths have nearly identical shape, and that our predicted dominant path lies almost directly on top of the mean path from the simulations.

	Rate of Escape	Mean Escape Time		
		$\epsilon = .05$	$\epsilon = .03$	$\epsilon = .015$
Baseline	3.14×10^{-4}	22.9	39.3	92.0
Lower Correlation	0.32	5.8	10.0	936.6
Tighter Intercept	0.28	167.3	803.1	-

TABLE 2. Rate of Escape (\bar{S}) and mean escape times for different specifications of the prior covariance matrix.

We then reconsider the cases in which we reduce the correlation between the slope and intercept, as in Figure 8 above, and when we tighten the prior on the intercept coefficient, as in Figure 9 above. Again we computed the dominant escape path for each specification, and then compared our findings to the results from 1000 simulations under different gain settings. In Figure 11 we plot the mean escape times for the different specifications. There we clearly see a large difference in the rate of escape. For each specification the escape times appear to increase roughly exponentially as the gain goes to zero, but the escape times increase fastest for the case where we tighten the prior on the intercept coefficient, and slowest in the baseline case. These results confirm our findings from solving the control problem, as Table 2 shows. We clearly see that the rate of escape (which determines the exponential rate of increase in the mean escape times) is orders of magnitude lower in the baseline specification than in the two alternatives, and the largest value comes with a tighter intercept prior. We've already seen in the simulations in Figures 8 and 9 above how the changes in prior beliefs alter the directions of escape, and in Figures 12 and 13 we plot the time paths of beliefs from representative samples. When the prior beliefs are that the slope and intercept of the Phillips curve are less highly correlated, as in Figure 12, then there is more variation in inflation outcomes prior to the final escape to zero inflation. The government attributes more of the variation in the data it observes as reflecting changes in the intercept coefficient, which in turn is associated with more uncertainty about the




FIGURE 11. Mean escape times from 1000 simulations for different prior specifications.




FIGURE 12. Mean escape paths and resulting inflation outcomes under a prior with lower correlation between the slope and intercept coefficients.

natural rate of unemployment. However with a tighter prior on the intercept coefficient, as in Figure 13, beliefs are tighter about the natural rate and there is less variation in inflation until it rapidly declines to zero. Thus different beliefs imply different escape paths, and therefore different time paths of inflation, although in each case the escapes lead to the Ramsey outcome of zero inflation.




FIGURE 13. Mean escape paths and resulting inflation outcomes under a tighter prior on the intercept coefficient.

8. SIMS'S NONCONVERGENCE EXAMPLE

Sims (1988) provides an example of apparent nonconvergence in this model. His results rely on a slight difference between his specification and ours, and the fact that he does not let the rate of variation in the parameters go to zero. Altering either of those differences leads to convergence.

For most of the paper we have not distinguished between the standard deviation σ of the regression error in (3.13) and the standard deviation σ_1 of the unemployment shocks in (2.2). This is because CWS assume that the two are equal, which makes sense because in a self-confirming equilibrium then (3.13) and (2.2) coincide. With $\sigma = \sigma_1$ the variation that policymakers observe in the unemployment rate is correctly decomposed into variation in the regressors and variation due to exogenous shocks. However the specifications that Sims (1988) studies assume that $\sigma \neq \sigma_1$. The misallocation of the observed variation leads to prolonged, perhaps indefinitely prolonged, episodes in which agents are away from the self-confirming equilibrium. An example is shown by the solid line in Figure 14, which plots the simulated inflation rate in a specification analogous to one of Sims's benchmark cases. (The parameterization is as in his Chart 5, which in our notation is $u = 6$, $\sigma_1 = 0.5$, $\sigma_2 = 0.4$, $V = [0.01, 0; 0, 0.01]$ and $\sigma = 0.1$.) We take this to be a benchmark with $\epsilon = 1$. As in Sims (1988), we ran the simulation for a total of 10,000 periods and did not observe convergence to the self-confirming equilibrium. The mean inflation rate over this long sample is 0.61, and the long simulation looks qualitatively similar to the smaller sample shown.

It is interesting, although perhaps not surprising that we do not obtain convergence for this specification. All of our results above have considered limits as $\epsilon \rightarrow 0$, which




FIGURE 14. Simulations of inflation from Sims’s specification, and some variations on it.

implies that we shrink the rate of prior variation in the parameters. Thus our results are approximations for small parameter drift, where “small” is context specific. Recall from (3.19) that the key term in the belief updating equation accounting for parameter variation is $\sigma^{-2}V = \epsilon^2\sigma^{-2}\hat{V}$. As we let $\epsilon \rightarrow 0$ we do get convergence to the SCE, as the dashed line in Figure 14 shows. There we retain the same specification as above, but now let $\epsilon = 0.2$, which effectively reduces the parameter variation by a factor of 0.04. Here we see that we rather quickly move toward the SCE, which is associated with an inflation rate of 6% in this parameterization. Exactly the same simulated time path results when we change the specification so that $\sigma = \sigma_1$ as we have assumed elsewhere. In Sims’s parameterization, σ_1 is five times as large as σ , so increasing σ by a factor of five is equivalent to reducing the gain ϵ to one fifth.

Thus we can interpret the difference between our results and those of Sims (1988) in two equivalent ways. First, we can say that the time variation allowed for in the parameters was simply too large, and that by shrinking it, convergence prevails. Alternatively, for the same amount of time variation in the parameters agents’ to do not attribute enough variation to the regression error. These interpretations lead to slightly different limit conclusions (i.e. of whether $\sigma = \sigma_1$ and the form of \hat{V}), but they lead to the same qualitative result: convergence to the self-confirming equilibrium.

9. CONCLUSION

This paper has reconciled and extended earlier work by Sims (1988), Chung (1990), Cho, Williams, and Sargent (2002), and Sargent (1999) that analyzed a model whose dynamics are driven by the decisions of an adaptive government that acts as a Bayesian in recursively estimating a model with drifting coefficients. The principal contributions of our paper are first, to extend the class of government priors by allowing an arbitrary covariance

matrix to govern the increments to coefficients in the government's model; and second, to apply recent results of Williams (2002) to the analysis of both the mean dynamics that govern convergence to a self-confirming equilibrium and the escape dynamics that govern recurrent excursions toward the vicinity of a Ramsey outcome. By altering the covariance matrix of increments to the government's coefficients, we have been able to study how tightening the prior in various directions affects both convergence and escape. We have also discovered that for some specifications of priors, the mean dynamics converge not to a self-confirming equilibrium but to a limit cycle that passes near the Ramsey outcome.

We believe that the concept of a self-confirming equilibrium is a good one for macroeconomics partly because it allows the meaningful disagreements about model structure that we witness in monetary and fiscal policy discussions, without surrendering the empirical discipline that is supplied by rational expectations. Self-confirming equilibrium is also appealing for reasons advocated by Fudenberg and Levine (1998), namely, its status as a limit point of a learning process in which agents confront their subjective models with the data they observe. Once we acknowledge that agents' models may be misspecified, we seem impelled to analyze the properties of self-confirming equilibria and the process of attaining them.

REFERENCES

- Arrow, K. J. and M. McManus (1958). A Note on Dynamic Stability. *Econometrica* 26, 448–454.
- Benveniste, A., M. Métivier, and P. Priouret (1990). *Adaptive Algorithms and Stochastic Approximations*. New York: Springer-Verlag.
- Bullard, J. (1994). Learning Equilibria. *Journal of Economic Theory* 64, 468–485.
- Cho, I.-K., N. Williams, and T. J. Sargent (2002). Escaping Nash Inflation. *Review of Economic Studies* 69, 1–40.
- Chung, H. (1990). *Did Policy Makers Really Believe in the Phillips Curve?: An Econometric Test*. Ph. D. thesis, University of Minnesota.
- Conley, T. G., L. P. Hansen, E. G. J. Luttmer, and J. A. Scheinkman (1997). Short-Term Interest Rates as Subordinated Diffusions. *Review of Financial Studies* 10, 525–577.
- Evans, G. W. and S. Honkapohja (2001). *Learning and Expectations in Macroeconomics*. Princeton University Press.
- Fudenberg, D. and D. Levine (1998). *The Theory of Learning in Games*. MIT. Press.
- Honkapohja, S. and K. Mitra (2002). Learning Stability in Economies with Heterogeneous Agents. Working paper, University of Helsinki.
- Kushner, H. J. and G. G. Yin (1997). *Stochastic Approximation Algorithms and Applications*. Springer-Verlag.
- Kydland, F. and E. C. Prescott (1977). Rules Rather than Discretion: the Inconsistency of Optimal Plans. *Journal of Political Economy* 85, 473–491.
- Mil'shtein, G. N. and L. B. Ryashko (1995). A First Approximation of the Quasipotential in Problems of the Stability of Systems with Random Non-degenerate Perturbations. *Journal of Applied Mathematics and Mechanics* 59(1), 47–56.
- Perko, L. (1996). *Differential Equations and Dynamical Systems*. New York: Springer-Verlag.

- Sargent, T. J. (1999). *The Conquest of American Inflation*. Princeton University Press.
- Sims, C. A. (1988). Projecting Policy Effects with Statistical Models. *Revista de Analisis Economico* 3, 3–20.
- Williams, N. (2001). *Escape Dynamics in Learning Models*. Ph. D. thesis, University of Chicago.
- Williams, N. (2002). *Escape Dynamics in Learning Models*. Working paper, Princeton University.

THOMAS J. SARGENT, DEPARTMENT OF ECONOMICS, NEW YORK UNIVERSITY, NEW YORK, NEW YORK

E-mail address: `thomas.sargent@nyu.edu`

NOAH WILLIAMS, DEPARTMENT OF ECONOMICS, PRINCETON UNIVERSITY, PRINCETON, NEW JERSEY

E-mail address: `noahw@princeton.edu`