Review of Economic Studies (2002) **69**, 1–40 0034-6527/02/00010001$02.00

Escaping Nash Inflation

IN-KOO CHO
University of Illinois

NOAH WILLIAMS
Princeton University

and

THOMAS J. SARGENT
Stanford University

First version received July 1999; *final version accepted June* 2001 (*Eds.*)

An ordinary differential equation (ODE) gives the *mean dynamics* that govern the convergence to self-confirming equilibria of self-referential systems under discounted least squares learning. Another ODE governs *escape dynamics* that recurrently propel away from a self-confirming equilibrium. In a model with a unique self-confirming equilibrium, the escape dynamics make the government discover too strong a version of the natural rate hypothesis. The escape route dynamics cause recurrent outcomes close to the Ramsey (commitment) inflation rate in a model with an adaptive government.

"If an unlikely event occurs, it is very likely to occur in the most likely way."
Michael Harrison

1. INTRODUCTION

Building on work by Sims (1988) and Chung (1990), Sargent (1999) showed that a government that adaptively fits an approximating Phillips curve model will recurrently escape from the suboptimal time-consistent (or Nash) inflation rate and for many periods will set inflation near the optimal time-inconsistent outcome. However, later the government's view of the world changes in ways that induce it to make inflation gradually return to the time-consistent suboptimal outcome of Kydland and Prescott (1977). The time consistent outcome is a self-confirming equilibrium and a limit point of the system under least squares learning. The superior outcomes during recurrent escapes from the time-consistent outcome emerge because the government temporarily learns an approximate version of the natural rate hypothesis. These temporary escapes from the time-consistent outcome reflect a remarkable type of escape dynamics from a self-confirming equilibrium. These dynamics promote experimentation and are induced by unusual shock patterns that interact with the government's adaptive algorithm and its imperfect model.

The escapes from Nash inflation lead to dramatic changes in the government's inflation policy as it temporarily overcomes its inflationary bias. Some simulated time paths of inflation for different specifications of the model are shown in Figure 1. Inflation starts and remains near the high time-consistent value for a while, is rapidly cut to zero, but then gradually approaches the time-consistent high value again. This paper explains the dynamic forces that drive these outcomes; the "mean dynamics" that govern the return to self-confirming equilibrium and the "escape dynamics" that expel the system away from the self-confirming equilibrium.

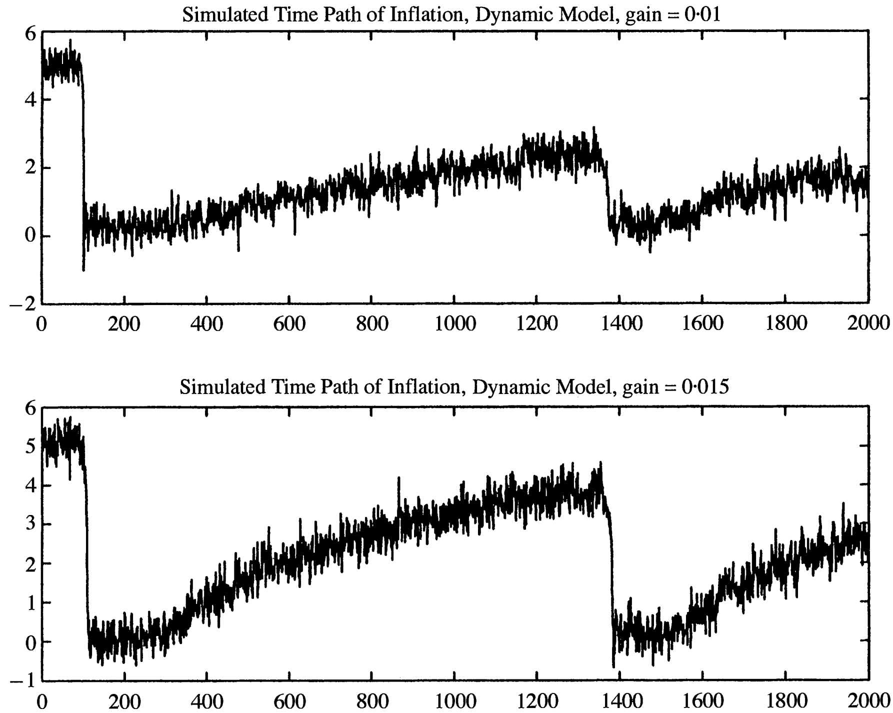

FIGURE 1

Simulated time paths of inflation for different specifications of the model

Escape dynamics from self-confirming equilibria can occur in a variety of models with large agents who use adaptive algorithms to estimate approximating models.[1] For concreteness, this paper focuses on the Phillips curve model studied by Sargent (1999). The model has the following features: (1) the monetary authority controls the inflation rate, apart from a random disturbance; (2) the true data generating mechanism embodies a version of the natural rate hypothesis in an expectational Phillips curve; (3) as in Kydland and Prescott (1977), a purposeful government dislikes inflation and unemployment and a private sector forecasts inflation optimally; but (4) the monetary policy makers do not know the true data generating mechanism and instead use a good fitting approximating model.[2] The fundamentals in the economy are fixed, including the true data generating mechanism, preferences, and agents' methods for constructing behaviour rules. Changes in the government's beliefs about the Phillips curve, and how it approximates the natural rate hypothesis, drive the inflation rate.

The self-confirming equilibrium concept is a natural one for analysing behaviour induced by an approximating model.[3] In a self-confirming equilibrium, beliefs are correct about events that occur with positive probability in equilibrium. The approximating model is "wrong" only in describing events that occur with zero probability in equilibrium.

1. See Williams (2001) and Bullard and Cho (2001) for some additional examples.

2. Inspired by econometric work about approximating models by Sims (1972) and White (1982), we endow the monetary authority, not with the correct model, but with an approximating model that it nevertheless estimates with good econometric procedures.

3. See Fudenberg and Levine (1993), Fudenberg and Kreps (1995), and Sargent (1999).

Among the objects determined by a self-confirming equilibrium are the parameters of the government's approximating model. While the self-confirming equilibrium concept differs formally from a Nash (or time consistent) equilibrium,[4] it turns out that the self-confirming equilibrium outcomes are the time-consistent ones. Thus, the suboptimal time consistent outcome continues to be our benchmark.

Like a Nash equilibrium, a self-confirming equilibrium restricts population objects (mathematical expectations, not sample moments). Our adaptive models are cast in terms of sample moments. We add adaptation by requiring the government to estimate its model from historical data in real time. We form an adaptive model by having the monetary authority adjust its behaviour rule in light of the latest model estimates. Thus, we attribute "anticipated utility" behaviour (see Kreps (1998)) to the monetary authority. Following Sims (1988), we study a "constant gain" estimation algorithm that discounts past observations. Called a "tracking algorithm", it is useful when parameter drift is suspected (see *e.g.* Marcet and Nicolini (1997)).

Results from the literature on least squares learning (*e.g.* Marcet and Sargent (1989*a*), Woodford (1990), Evans and Honkapohja (1998)) apply and take us part way, but only part way, to our goal of characterizing the dynamics of the adaptive system. That literature shows how an ordinary differential equation called the "mean dynamics" describes the limiting behaviour of systems with least squares learning. The mean dynamics describe the (unconditionally) average path of the government's beliefs, in a sense that we shall describe precisely. For our model, the mean dynamics converge to the self-confirming equilibrium and the time consistent outcome. Thus, the mean dynamics do not account for the recurrent stabilizations in the simulations of Sims (1988), Chung (1990), and Sargent (1999). We show that these stabilizations are governed by another *deterministic* component of the dynamics, described by another ODE, the "escape" dynamics. They point away from the self-confirming equilibrium and toward the Ramsey (or optimal-under-commitment) equilibrium outcome. So two sorts of dynamics dominate the behaviour of the adaptive system.

1. The *mean dynamics* come from an unconditional moment condition, the least squares normal equations. These dynamics drive the system *toward* a self-confirming equilibrium.[5]
2. The *escape route dynamics* propel the system *away* from a self-confirming equilibrium. They emerge from the same least squares moment conditions, but they are *conditioned* on a particular "most likely" unusual event, defined in terms of the disturbance sequence. This most likely unusual event is endogenous.

The escape route dynamics have a compelling behavioural interpretation. Within the confines of its approximate model, learning the natural rate hypothesis requires that the government generate a sufficiently wide range of inflation experiments. To learn even an imperfect version of the natural rate hypothesis, the government must experiment more than it does within a self-confirming equilibrium. The government is caught in an experimentation trap. The adaptive algorithm occasionally puts enough movement into the government's beliefs to produce informative experiments.

4. It is defined in terms of different objects.

5. But as Evans and Honkapohja (2001) point out, the route can be circuitous. In Figures 8 and 9, we indicate how for our model the mean dynamics point in the same direction as the escape dynamics along much of the escape route.

1.1. *Related literature*

Evans and Honkapohja (1993) investigated a model with multiple self-confirming equilibria having different rates of inflation. When agents learn through a recursive least squares algorithm, outcomes converge to a self-confirming equilibrium that is stable under the learning algorithm. When agents use a fixed gain algorithm, Evans and Honkapohja (1993) demonstrated that the outcome oscillates among different locally stable self-confirming equilibria. They suggested that such a model can explain wide fluctuations of market outcomes in response to small shocks.

In models like Evans and Honkapohja (1993) and Kasa (1999), the time spent in a neighbourhood of a locally stable equilibrium and the escape path from its basin of attraction are determined by a large deviation property of the recursive learning algorithm. As the stochastic perturbation disappears, the outcome stays in a neighbourhood of a particular locally stable self-confirming equilibrium (exponentially) longer than the others. This observation provided Kandori, Mailath and Rob (1993) and Young (1993) with a way to select a unique equilibrium in evolutionary models with multiple locally stable Nash equilibria.

An important difference from the preceding literature is that our model has a unique self-confirming equilibrium. Despite that, the dynamics of the model resemble those for models with multiple equilibria such as Evans and Honkapohja (1993). With multiple locally stable equilibria, outcomes escape from the basin of attraction of a locally stable outcome to the neighbourhood of another locally stable equilibrium. The fact that our model has a globally unique stable equilibrium creates an additional challenge for us, namely, to characterize the most likely direction of the escape from a neighbourhood of the unique self-confirming equilibrium. As we shall see, the most likely direction entails the government's learning a good, but not self-confirming, approximation to the natural rate hypothesis.

1.2. *Organization*

Section 2 describes the model in detail. Section 3 defines a self-confirming equilibrium. Section 4 describes a minimal modification of a self-confirming equilibrium formed by giving the government an adaptive algorithm for its beliefs. Section 5 uses results from the theory of large deviations developed in Williams (2001) to characterize convergence to and escape from a self-confirming equilibrium. Section 6 shows that numerical simulations of escape dynamics, like those in Sargent (1999), are well described by the numerically calculated theoretical escape paths. For the purpose of giving intuition about the escape dynamics, Section 7 specializes the shocks to be binomial, then adduces a transformed measure of the shocks that tells how particular endogenously determined unusual shock sequences drive the escape dynamics. Section 8 concludes. The remainder of this introduction describes the formal structure of the model and findings of the paper.

1.3. *Overview*

The government's beliefs about the economy are described by a vector of regression coefficients γ. It chooses a decision rule $h(\gamma)$ that makes the stochastic process ξ for the economy be $\xi(\gamma)$. But for the stochastic process $\xi(\gamma)$, the best fitting model of the economy has coefficients $\Gamma = T(\gamma)$. A self-confirming equilibrium is a fixed point of $T(\gamma)$. The orthogonality conditions pinning down the best fitting model can be expressed

$$Eg(\gamma, \xi) \equiv \bar{g}(\gamma) = 0, \tag{1.1}$$

where E is the mathematical expectation with respect to the distribution of $\xi(\gamma)$. We shall show that

$$\bar{g}(\gamma) = \bar{M}(T(\gamma) - \gamma),$$

where $\bar{M}$ is the second moment matrix of the right side variables in the government's model. Thus, a self-confirming equilibrium solves

$$\bar{g}(\gamma) = 0. \tag{1.2}$$

A self-confirming equilibrium is a set of population regression coefficients. We form an adaptive model by slightly modifying a self-confirming equilibrium. Rather than using population moments to fit its regression model, the government uses discounted least squares estimates from historical samples. We study how the resulting adaptive system converges to or diverges from a self-confirming equilibrium. Each period the government uses the most recent data to update a least squares estimate γ_n of its model coefficients γ, then sets its policy according to $h(\gamma_n)$. This is what Kreps (1998) calls an anticipated utility model. The literature on least squares learning in self-referential systems (see Marcet and Sargent (1989*a*), Marcet and Sargent (1989*b*), Woodford (1990), and Evans and Honkapohja (1998)) gives conditions under which the limiting behaviour of the government's beliefs are nearly deterministic and approximated by the following ordinary differential equation (ODE):

$$\dot{\gamma} = R^{-1}\bar{g}(\gamma) \tag{1.3}$$

$$\dot{R} = M - R. \tag{1.4}$$

Equations (1.3), (1.4) define the *mean dynamics*. A fixed point $\bar{\gamma}$ of the ODE (1.3), (1.4) is a self-confirming equilibrium ($\bar{g}(\bar{\gamma}) = 0$ with $R = M$). The least squares learning literature describes how the convergence of γ_n to $\bar{\gamma}$ is governed by the uniqueness and stability of the stationary points of the ODE.

Our model has a unique self-confirming equilibrium. It supports the high inflation time-consistent outcome of Kydland and Prescott (1977). The ODE (1.3), (1.4), is very informative about the behaviour of our adaptive model. It is globally stable about the self-confirming equilibrium, and describes how the adaptive system is gradually drawn to the self-confirming equilibrium.[6]

But to understand how the sample paths recurrently visit the better low-inflation outcome, we need more than the ODE (1.3), (1.4).

6. Let $\dot{x} = \bar{b}(x)$ be an ordinary differential equation with stationary solution $x = x^*$. We say that x^* is locally stable if there exists an open neighbourhood Λ of x^* such that for any compact set $G \subset \Lambda$ and any $\delta > 0$, there exists $\tau^* < \infty$ such that all trajectories of the ordinary differential equations originating in G are in a δ neighbourhood of x^* for all $\tau > \tau^*$. If the same condition holds for any open neighbourhood of x^*, then we say that x^* is globally stable. (*cf.* Dupuis and Kushner (1989).)

Before now, "escape dynamics" have been simulated and their causes described informally (see Sargent (1999)), but they have not been completely characterized analytically. This paper shows that they are governed by an ODE of the form

$$\dot{\gamma} = R^{-1}\bar{g}(\gamma) + \dot{v} \tag{1.5}$$

$$\dot{R} = M - R \tag{1.6}$$

$$\dot{v} = v(\gamma, R). \tag{1.7}$$

We display a model approximation problem whose solution generates (1.5) and use it to interpret $\dot{v}$ as a continuous time limit of the orthogonality conditions (1.1) under a twisted distribution for the shock process. This twisted distribution is the "most likely unlikely" shock process. Twisting the orthogonality conditions results in "endogenous experimentation" that makes the government learn an approximate version of the natural rate hypothesis. Thus, like the mean dynamics, the escape dynamics are deterministic. We verify that these deterministic dynamics do a good job of describing the simulations.

As Sims (1988) and Sargent (1999) emphasize, the evolution of beliefs during an escape is economically interesting. During escapes, the government discovers a good enough approximate version of the natural rate hypothesis to cause it to pursue superior policy. The policy is supported by beliefs that are "wrong" in the sense that they are not a self-confirming equilibrium. Nevertheless, in another sense those beliefs are more "correct" than those in a self-confirming equilibrium because they inspire the government to leave the "experimentation trap" that confines it within a self-confirming equilibrium.

2. SETUP

In the model, time is discrete and indexed by n. Let $W'_n = [W_{1n} \quad W_{2n}]$ be an i.i.d. sequence of (2×1) random vectors with mean zero and covariance matrix I. Let $U, \pi, \hat{x}, x$, respectively, be the unemployment rate, the rate of inflation, the public's expected rate of inflation, and the systematic part of inflation determined by government policy. The government sets x, the public sets $\hat{x}$, then nature chooses shocks W that determine π and U. The economy is described by the following version of a model of Kydland and Prescott (1977):

$$U_n = u - \theta(\pi_n - \hat{x}_n) + \sigma_1 W_{1n}, \; u > 0, \; \theta > 0 \tag{2.8}$$

$$\pi_n = x_n + \sigma_2 W_{2n} \tag{2.9}$$

$$\hat{x}_n = x_n \tag{2.10}$$

$$x_n = h(\gamma)' X_{n-1} \tag{2.11}$$

where

$$X_{n-1} = [U_{n-1} \quad U_{n-2} \quad \pi_{n-1} \quad \pi_{n-2} \quad 1]'. \tag{2.12}$$

Equation (2.8) is a natural rate Phillips curve; (2.9) says that the government sets inflation up to a random term; (2.10) imposes rational expectations for the public; (2.11) is the government's decision rule for setting the systematic part of inflation x_n. The decision rule $h(\gamma)$ is a function of the government's beliefs about the economy, which are parameterized by a vector γ.

2.1. *The government's beliefs and control problem*

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

$$U_n = \gamma_1 \pi_n + \gamma'_{-1} X_{n-1} + \eta_n, \tag{2.13}$$

where the government treats η as a mean zero, serially uncorrelated random term beyond its control.[7] We shall eventually restrict γ, but temporarily regard it as arbitrary. The government's decision rule (2.11) solves the problem:

$$\min_{\{x_n\}} \hat{E} \sum_{n=0}^{\infty} \delta^n (U_n^2 + \pi_n^2) \tag{2.14}$$

where $\hat{E}$ denotes the expectations operator induced by (2.13) and the minimization is subject to (2.13) and (2.9).

We call problem (2.14) the *Phelps problem*. Versions of it were studied by Phelps (1967), Kydland and Prescott (1977), Barro and Gordon (1983), and Sargent (1999). We identify three salient outcomes associated with different hypothetical government's beliefs:

- *Belief* 1. If $\gamma_1 = -\theta, \gamma_{-1} = [0 \quad 0 \quad 0 \quad 0 \quad u]$, then the Phelps problem tells the government to set $x_n = \theta u$ 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 $\gamma_1 = 0$, $\gamma_{-1} = [0 \quad 0 \quad 0 \quad 0 \quad u^*]$ 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 $\delta \to 1$ from below, the Phelps problem eventually sends x_n arbitrarily close to 0.

Under the actual probability distribution generated by (2.8), (2.9), (2.10), the value of the government's objective function (2.14) is larger under the outcome $x_n = 0$ than under $x_n = \theta u > 0$. Under Belief 1, the government perceives a trade-off between inflation and unemployment and sets inflation above zero to exploit that trade-off. 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 although there is a short-term trade-off between inflation and unemployment when $\gamma_1 < 0$, there is no "long-term" trade-off. An "induction hypothesis" opens an avenue by which the government can manipulate the future positions of the Phillips curve (see Cho and Matsui (1995) and Sargent (1999)). The Phelps problem then tells the government eventually to set inflation close to zero when δ is close to 1.

In a common-knowledge model in which (2.13) is dropped and replaced by the assumption that the government knows the model, the outcome $x_n = u\theta$ emerges as what Stokey (1989) and Sargent (1999) call the Nash outcome, and $x_n = 0$ emerges as the Ramsey outcome. In the common-knowledge model, these varying outcomes reflect different timing protocols and characterize a time-consistency problem analysed by Kydland and Prescott. The mapping from government beliefs to outcomes is interesting only when the government's beliefs might be free. Our equilibrium concept, a self-confirming equi-

7. For 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 (2.8), (2.9), all of the temporal dependence in the model comes through the government's beliefs. For the static model, the government's control rule can be calculated explicitly, allowing some of our characterizations to be sharper.

librium, restricts those beliefs, and thereby narrows the outcomes relative to those enumerated above. However, the mapping from beliefs to outcomes play a role during escapes from self-confirming equilibria.

3. SELF-CONFIRMING EQUILIBRIUM

3.1. *Restrictions on government's beliefs*

Define $\xi_n = [W_{1n} \quad W_{2n} \quad X'_{n-1}]'$ and

$$g(\gamma, \xi_n) = \eta_n \begin{bmatrix} \pi_n \\ X_{n-1} \end{bmatrix}. \tag{3.15}$$

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:

$$0 = E\left(\eta_n \begin{bmatrix} \pi_n \\ X_{n-1} \end{bmatrix} \right). \tag{3.16}$$

Equations (3.16) are the orthogonality conditions that make γ in (2.13) a least-squares regression. Condition (3.18) thus renders the government's beliefs consistent with the data.

Let W^n denote the history of the joint shock process $[{}^{W_{1n}}_{W_{2n}}]$ up to n. Evidently, from (2.8), (2.9), (2.10), (2.11), X_{n-1} and therefore the ξ_n process are both functions of γ:

$$\xi_n = \xi(\gamma, W^n). \tag{3.17}$$

Definition 3.1. *A self-confirming equilibrium is a* γ *that satisfies*

$$Eg(\gamma, \xi_n) = 0. \tag{3.18}$$

The expectation in (3.18) is taken with respect to the probability distribution generated by (2.8), (2.9), (2.10), (2.11).

Condition (3.18) can be interpreted as asserting that γ is a fixed point in a mapping from the government's beliefs about the Phillips curve to the actual Phillips curve. Thus, let

$$M_n = \begin{bmatrix} \pi_n \\ X_{n-1} \end{bmatrix} \begin{bmatrix} \pi_n \\ X_{n-1} \end{bmatrix}', \qquad \bar{M} = EM_n. \tag{3.19}$$

Let $\Gamma = T(\gamma)$ be the least squares regression coefficients in $U_n = \Gamma'[{}^{\pi_n}_{X_{n-1}}] + \nu_n$ where ν_n is a least squares residual orthogonal to the regressors. We write $\Gamma = T(\gamma)$ because via the government best response mapping $h(\gamma)$, Γ depends on γ through the moment matrices $\bar{M}$ and $EU_n[{}^{\pi_n}_{U_{n-1}}]$. Then notice that

$$Eg(\gamma, \xi_n) = E\begin{bmatrix} \pi_n \\ X_{n-1} \end{bmatrix}(U_n - [\pi_n \quad X'_{n-1}]\gamma) \tag{3.20}$$

$$= E\left(U_n \begin{bmatrix} \pi_n \\ X_{n-1} \end{bmatrix} - \bar{M}\gamma \right) \tag{3.21}$$

$$= \bar{M}(\Gamma - \gamma). \tag{3.22}$$

Given a government model in the form of a perceived regression coefficient vector γ and the associated government best response function $h(\gamma)$, $\Gamma = \bar{M}^{-1}EU_n[{}^{\pi_n}_{X_{n-1}}] = T(\gamma)$ is the actual least squares regression coefficient induced by $h(\gamma)$. Thus, T maps government model γ to a best fitting model $T(\gamma)$. Equation (3.22) shows that (3.18) asserts that $T(\gamma) = \gamma$, so that the government's model is the best fitting model. See Marcet and Sargent (1989*a*) for a discussion of the T operator in a related class of models.

Elementary calculations show that there is a unique self-confirming equilibrium. It corresponds to Belief 1 mentioned above and supports the Nash equilibrium outcome in the sense of Stokey (1989) and Sargent (1999).

4. ADAPTATION

4.1. *Discounted least squares updating of γ*

We modify the model now to consist of (2.8), (2.9), (2.10) as before, but replace (2.11) with

$$x_n = h(\gamma_n)' X_{n-1} \tag{4.23}$$

where $h(\gamma)$ remains the best-response function generated by the Phelps problem, and γ_n is the government's time n estimate of the empirical Phillips curve. The government estimates γ by the following recursive least squares algorithm:

$$\gamma_{n+1} = \gamma_n + \varepsilon R_n^{-1} g(\gamma_n, \xi_n) \tag{4.24}$$

$$R_{n+1} = R_n + \varepsilon(M_n - R_n) \tag{4.25}$$

where ε is a gain parameter that determines the weight placed on current observations relative to the past. In this paper we consider the case in which the gain is constant. We want to study the behaviour of system formed by (2.8), (2.9), (2.10), (4.23), (4.24) and (4.25).

4.2. *Mean dynamics*

We find the first important component of dynamics by adapting the stochastic approximation methods used by Woodford (1990), Marcet and Sargent (1989*a*), and Evans and Honkapohja (2001). We call this component the *mean dynamics* because it governs the (unconditionally) expected evolution of the government's beliefs. While previous applications of stochastic approximation results in economics have generally considered recursive least squares with decreasing gain, we consider the case where the gain is constant.[8] Broadly similar results obtain in the constant and decreasing gain cases, but there are important differences in the asymptotics and the sense of convergence that we discuss below.

To present convergence proofs, it helps to group together the components of the government's beliefs into a single vector. Define

$$\theta_n = \begin{bmatrix} \gamma_n \\ \text{col}\,(R_n) \end{bmatrix}, \qquad Z_n = \begin{bmatrix} R_n^{-1} g(\gamma_n, \xi_n) \\ \text{col}\,(M_n - R_n) \end{bmatrix}. \tag{4.26}$$

8. See Evans and Honkapohja (2001) for extensive discussion of constant gain algorithms.

Then the updating equations (4.24), (4.25) can be written

$$\theta_{n+1} = \theta_n + \varepsilon Z_n. \tag{4.27}$$

Now break the "update part" Z_n into its expected and random components. Define $v_n = Z_n - \bar{b}(\theta_n)$ where $\bar{b}$ is the mean of Z_n defined as

$$\bar{b}(\theta) = EZ_n = \begin{bmatrix} R^{-1}\bar{g}(\gamma) \\ \text{col}\,(\bar{M}(\gamma) - R) \end{bmatrix} \tag{4.28}$$

where

$$\bar{g}(\gamma) = Eg(\gamma, \xi), \qquad \bar{M}(\gamma) = EM_n. \tag{4.29}$$

Then we can write the composite dynamics as

$$\theta_{n+1} = \theta_n + \varepsilon\bar{b}(\theta_n) + \varepsilon v_n \tag{4.30}$$

To determine the expected evolution of the government's estimates, we study the asymptotic behaviour of the difference equation (4.30). Our convergence theorem is about sequences of economies where the gain goes to zero. This differs from typical applications of least squares learning, in which the gain sequence decreases (usually as $1/n$) over time. As in the decreasing gain case, we can show that the asymptotic behaviour of (4.30) is governed by an ODE, but the estimates converge in a weaker sense. Specifically, decreasing gain algorithms typically converge with probability one *along* a sequence of iterations as $n \to \infty$, but constant gain algorithms converge weakly (or in distribution) as $\varepsilon \to 0$ *across* sequences of iterations, each of which is indexed by the gain.

Note that we can rewrite (4.30) as

$$\frac{\theta_{n+1} - \theta_n}{\varepsilon} = \bar{b}(\theta_n) + v_n. \tag{4.31}$$

This equation resembles a finite-difference approximation of a derivative with time step ε, but is perturbed by a noise term. The convergence argument defines a continuous time scale as $t = n\varepsilon$, and interpolates between the discrete iterations to get a continuous process. Then by letting $\varepsilon \to 0$, the approximation error in the finite difference goes to zero, and a weak law of large numbers insures that the noise term v_n becomes negligible. We are left with the ODE:

$$\dot{\gamma} = R^{-1}\bar{g}(\gamma) \tag{4.32}$$

$$\dot{R} = \bar{M}(\gamma) - R. \tag{4.33}$$

We need the following set of assumptions. For reference, we also list the original number in Kushner and Yin (1997). To emphasize the asymptotics, we include the superscript ε on the parameters θ_n^ε denoting the gain setting.

Assumptions A.

A8.5.0. The random sequence $\{\theta_n^\varepsilon; \varepsilon, n\}$ is tight.[9]
A8.5.1. For each compact set D, $\{Z_n^\varepsilon 1_{\{\theta_n^\varepsilon \in D\}}; \varepsilon, n\}$ is uniformly integrable.[10]
A8.5.3. For each compact set D, the sequence $\{\bar{b}(\theta_n^\varepsilon)1_{\{\theta_n^\varepsilon \in D\}}; \varepsilon, n\}$ is uniformly integrable.
A8.5.4a. The ODE $\dot{\theta} = \bar{b}(\theta)$ has a point $\bar{\theta}$ that is asymptotically stable.[11]
A8.1.6. The function $\bar{b}(\theta)$ is continuous.
A8.1.7. For each $\delta > 0$, there is a compact set D_δ such that $\inf_{n,\varepsilon} P(\upsilon_n^\varepsilon \in D_\delta) \geqq 1 - \delta$.

The following theorem is based on results in Kushner and Yin (1997).

Theorem 4.1. *Under Assumptions* A*, as* $\varepsilon \to 0$ *the parameter sequence* θ_n^ε *converges weakly to the process* $\theta(t)$ *that solves the ordinary differential equations* (4.32, 4.33). *Further, Assumptions* A *hold for our model when the shocks* W_n *are i.i.d. normal.*

Proof. See Appendix A. ||

The theorem shows that the trajectories of the estimates converge to the trajectory of the ODE system. Since the ODE has a unique stable point $\bar{\theta}$ that is the self-confirming equilibrium, the estimate sequence converges weakly to the self-confirming equilibrium. Therefore, with high probability, as $\varepsilon \to 0$ and $n \to \infty$ we would expect the government's beliefs to be near their self-confirming values, and the economy to be near the Nash outcome. However, in the next section we shall see that the beliefs can recurrently escape the self-confirming equilibrium. Although the impact of noise terms goes to zero with the gain, for a *given* positive ε, "rare" sequences of shocks can have a large impact on the estimates and the economy.

5. ESCAPE

In this section we determine the most likely rare events and how they push the government's beliefs away from a self-confirming equilibrium. To this end, we first present some general results from the theory of large deviations, a general method for analysing small probability events.

5.1. *Escape dynamics as a control problem*

Throughout, we will only be interested in characterizing the escape problem for the Phillips curve coefficients γ. This motivates the following definition.

9. A random sequence $\{A_n\}$ is *tight* if

$$\lim_{K\to\infty} \sup_n P(|A_n| \geqq K) = 0.$$

10. A random sequence $\{A_n\}$ is *uniformly integrable* if

$$\lim_{K\to\infty} \sup_n E(|A_n| 1_{\{|A_n|\geqq K\}}) = 0.$$

11. $\bar{x}$ is *asymptotically stable* for an ODE if any solution $x(t) \to \bar{x}$ as $t \to \infty$, and for each $\delta > 0$ there exists an $\epsilon > 0$ such that if $|x(0) - \bar{x}| \leqq \epsilon$, then $|x(t) - \bar{x}| \leqq \delta$ for all t.

Definition 5.1. *An escape path is a sequence of estimates solving* (4.24) *that leave a set* G *containing the limit point* $\bar{\gamma}$:

$$\{\gamma_n\}_{n=0}^{N}, \gamma_0 \in G(\bar{\gamma} \in G), \gamma_n \notin G \qquad \text{for some } n \leqq N < \infty.$$

Following a convention in the large deviation literature, we set the initial point of an escape path to be the stable point $\bar{\gamma}$. Given a gain $\varepsilon > 0$ and a compact neighbourhood G of the stable point $\bar{\gamma}$, let $B^{\varepsilon}(G)$ be the set of all escape paths. For each $\{\gamma_n\} \in B^{\varepsilon}(G)$, define

$$\tau^{\varepsilon} = \varepsilon \inf\{n : \gamma_0 = \bar{\gamma}, \gamma_n \notin G\}$$

as the (first) exit time out of G. Because the exit time varies across different escape paths, $B^{\varepsilon}(G)$ induces a probability distribution over the exit times. We want to understand the probability distribution over $B^{\varepsilon}(G)$ and the probability distribution of exit times when $\varepsilon > 0$ is sufficiently small. In particular, we want to identify the most likely escape paths in $B^{\varepsilon}(G)$. Let ∂G be the boundary of G.

Definition 5.2. *Let* $\tau^{\varepsilon}(\{\gamma_n\})$ *be the* (*first*) *exit time associated with escape path* $\{\gamma_n\} \in B^{\varepsilon}(G)$. *An absolutely continuous trajectory* φ *is a* dominant escape path *if* $\varphi(0) = \bar{\gamma}$, $\varphi(\tau^*) \in \partial G$, *and for all* $\rho > 0$:

$$\lim_{\varepsilon \to 0} \Pr\left(\forall n < \tau^{\varepsilon}(\{\gamma_n\})/\varepsilon, \quad \exists \tau' \leqq \tau^*, |\varphi(\tau') - \gamma_n| < \rho : B^{\varepsilon}(G)\right) = 1. \tag{5.34}$$

Roughly speaking, (5.34) states that as the gain $\varepsilon > 0$ converges to 0, the set of escape paths converges to a small neighbourhood of φ. An escape from $\bar{\gamma}$ will occur along φ with very high probability, if an escape ever occurs.

To analyse the escape dynamics, we adapt the general results of Dupuis and Kushner (1989), which are themselves extensions of the theory of Freidlin and Wentzell (1984) to stochastic approximation models. After presenting some general results, we apply results of Williams (2001), who obtains explicit characterizations of the escape dynamics in a class of models. In our setting, these results can be used to interpret the simulations calculated earlier by Sims (1988), Chung (1990), and Sargent (1999). Given the recursive formula (4.30), for a vector α, define the H-functional as:

$$H(\gamma, \alpha; R) = \lim_{n \to \infty} \frac{1}{n} \log E_{\xi_0} \exp\left\langle \alpha, \textstyle\sum_{i=1}^{n} R^{-1} g(\gamma, \xi_i)\right\rangle \tag{5.35}$$

where $\langle \cdot, \cdot \rangle$ is the inner product of two vectors, and the expectation is conditioned on an arbitrary initial state ξ_0. This function averages over the time dependence in the shocks ξ_n to determine the large deviation properties of the γ_n parameter sequence. Note that it depends implicitly on the matrix R. We then define the Legendre transform of H as:

$$L(\gamma, \beta; R) = \sup_{\alpha} [\langle \alpha, \beta \rangle - H(\gamma, \alpha; R)]. \tag{5.36}$$

The *action functional* is defined over absolutely continuous trajectories $\varphi = (\gamma(t))_0^T$ by:

$$S(T, \varphi) = \int_0^T L(\gamma, \dot{\gamma}; R) ds \tag{5.37}$$

with $\gamma(0) = \bar{\gamma}$, and with the evolution of R following the mean dynamics conditional on γ. (We let $S = +\infty$ for trajectories that are not absolutely continuous.) In the context of continuous time diffusions, Freidlin and Wentzell (1984) characterized the dominant escape path as a solution of a variational problem. Their results have been extended to discrete time stochastic approximation models by Dupuis and Kushner (1985) and Dupuis and Kushner (1989). We adapt these results in the following theorem, whose main object is the solution of the following variational problem:

$$\bar{S} = \min_{\varphi} S(t, \varphi) \tag{5.38}$$

subject to

$$\dot{R} = \bar{M}(\gamma) - R$$

$$\gamma(0) = \bar{\gamma}, \qquad R(0) = \bar{R}, \qquad \gamma(t) \notin G \quad \text{for some } 0 < t < T.$$

The minimized value $\bar{S}$ is the *rate function* that determines the (exponential) bound for the large deviation estimates. The following theorem compiles and applies results from Dupuis and Kushner (1989), Kushner and Yin (1997), and Dembo and Zeitouni (1998). In the following theorem, we let $\gamma^{\varepsilon}(t)$ be the piecewise linear interpolation of the estimate sequence.

Theorem 5.3. *Suppose that Assumptions* A *hold.*

1. *Suppose that the shocks W are i.i.d. and unbounded but that there exists a σ−algebra $\mathcal{F}_n \supset \sigma(\gamma_i, i \leqq n)$ and constants $\kappa > 1, B < \infty$ such that for all n and $s \geqq 0$,*

$$P(|R_n^{-1} g(\gamma_n, \xi_n)| \geqq s | \mathcal{F}_n) \leqq B \exp(-s^{\kappa}) \ a.s.$$

Then we have:

$$\limsup_{\varepsilon \to 0} \varepsilon \log P(\gamma^{\varepsilon}(t) \notin G \quad \textit{for some } \ 0 < t \leqq T | \gamma^{\varepsilon}(0) = \bar{\gamma}) \leqq -\bar{S}.$$

2. *If the shocks W are i.i.d. and bounded, and $\bar{S}$ is continuous as a function of the radius of the set G then we have:*

$$\lim_{\varepsilon \to 0} \varepsilon \log P(\gamma^{\varepsilon}(t) \notin G \quad \textit{for some } 0 < t \leqq T | \gamma^{\varepsilon}(0) = \bar{\gamma}) = -\bar{S}.$$

3. *Under the assumptions of part* 2, *for any escape path with gain ε, let τ^{ε} be the time of first escape from G. Then for all $\delta > 0$:*

$$\lim_{\varepsilon \to 0} P\big(\exp\big((\bar{S} + \delta)/\varepsilon\big) > \tau^{\varepsilon} > \exp\big((\bar{S} - \delta)/\varepsilon\big)\big) = 1.$$

and

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

4. *Under the assumptions of part* 2, *let $\bar{x} = \gamma(\tau^{\varepsilon})$ be the terminal point of the dominant escape path. Then for any $\gamma^{\varepsilon}(\tau^{\varepsilon})$ and $\delta > 0$:*

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

Proof. See Appendix B. ||

This theorem establishes the precise sense in which the solution of the control problem determines the most likely escape path. In particular, part (1) shows that the probability of observing an escape episode is exponentially decreasing in the gain, with the rate given by the minimized value of the cost function $\bar{S}$. The next three parts establish stronger results under the assumption that the errors are bounded. Part (2) shows that under bounded errors, the asymptotic inequality in part (1) becomes an asymptotic equality. Part (3) shows that for small ε the time it takes beliefs to escape the self-confirming equilibrium becomes close to $\exp(\bar{S}/\varepsilon)$. Finally, part (4) shows that with probability approaching one, if the beliefs escape, they escape along the dominant escape path.

5.2. *Characterizing the escape dynamics*

While Theorem 5.3 offers a characterization of the dominant escape path, it is difficult to derive useful insights from the minimization problem itself. Additionally, because of the complicated nature of H and S, analysis of the escape dynamics and determination of the exponential rate of convergence appear to be daunting tasks. However Williams (2001) draws on the recent results of Worms (1999) to simplify the problem and to provide both an analytic characterization of the escape dynamics and a numerical solution. The key step is to recognize that although (5.35) is a complicated function, we can use some additional results from the applied probability literature to simplify the analysis.

In particular, Varadhan's theorem (see Dembo and Zeitouni (1998)) shows the duality between moments of exponential functions and large deviations. Since the H-functional in (5.35) is an asymptotic exponential moment, if we can identify a large deviation rate function for the $g(\gamma, \xi)$ process, we can identify H. The large deviation results of Worms (1999) identify this rate function in terms of the solution of the Poisson equation associated with $g(\gamma, \xi)$. It is known (see Benveniste, Metivier and Priouret (1990) for example) that the asymptotic distribution of Markov processes can be characterized by the Poisson equation, so it is natural that it appears here. This analysis then leads to a representation of the H-functional as a quadratic form in the vector α, with a normalizing matrix Q that depends on the solution of the Poisson equation associated with $g(\gamma, \xi)$. In general solving the Poisson equation can be difficult because it is a functional equation. However in the important linear-quadratic-Gaussian case (which includes our model), the problem can be solved in the space of quadratic functions. Therefore the Poisson equation reduces to some matrix Lyapunov equations. This provides a tremendous simplification, as there are efficient numerical methods for solving Lyapunov equations. We summarize these arguments in the following theorem and remark.

Theorem 5.4. *Suppose that Assumptions* A *hold, that ξ_n follows a stationary functional autoregression with a unique stationary distribution and Lipschitz continuous mean and variance functions, and that the function $g(\gamma, \xi)$ is Lipschitz continuous in ξ. Then there is a matrix-valued function $Q(\gamma, R)$ such that the dominant escape path and rate function can be determined by solving the following variational problem:*

$$\bar{S} = \inf_{\dot{v}} \frac{1}{2} \int_0^t \dot{v}(s)' Q(\gamma(s), R(s))^{-1} \dot{v}(s) ds \tag{5.39}$$

subject to

$$\dot{\gamma} = R^{-1} \bar{g}(\gamma) + \dot{v} \tag{5.40}$$

$$\dot{R} = \bar{M}(\gamma) - R \tag{5.41}$$

$$\gamma(0) = \bar{\gamma},\ R(0) = \bar{R},\ \gamma(t) \notin G \qquad \textit{for some } 0 < t < T. \tag{5.42}$$

Proof. See Williams (2001). ||

Remark 5.5. *In our model, ξ_n follows a linear autoregression, the W_{in} are i.i.d. normal, and $g(\gamma, \xi)$ is a quadratic function of ξ. Then $Q(\gamma, R)$ is a fourth moment matrix that can be calculated explicitly by solving the matrix Lyapunov equations described in Appendix C.*

This theorem provides a clearer interpretation and analysis of the variational problem. The escape dynamics perturb the mean dynamics by a forcing sequence $\dot{v}$. Then S is a quadratic cost function that measures the magnitude of the perturbations during the episode of an escape. In particular, we can think of (5.39) as a least squares problem, where Q plays the role of a covariance matrix. If we had $\dot{v} \equiv 0$ then the beliefs adhere to the mean dynamics, and the cost would be zero. For the beliefs to escape from $\bar{\gamma}$, we require nonzero perturbations. To find the most likely escape path we want to perturb the evolution as little as possible. In other words, it takes a sequence of unusual events to push γ away from the self-confirming equilibrium. The control problem (5.39) says that to find the dominant escape path, we need to look for a least cost sequence of shocks that will push beliefs away from $\bar{\gamma}$.

To find the dominant escape path, we solve the control problem in (5.39). We form the Hamiltonian with co-states (a, λ) for the evolution of (γ, R):

$$\mathcal{H} = a \cdot R^{-1}\bar{g}(\gamma) + \tfrac{1}{2}a'Q(\gamma, R)a - \lambda \cdot \bar{M}(\gamma).$$

It is easy to verify that the Hamiltonian is convex, so the first-order conditions are necessary and sufficient for a minimum. Taking the first-order conditions, we see that the dominant escape path is characterized by the following set of differential equations:

$$\begin{aligned}
\dot{\gamma} &= R^{-1}\bar{g}(\gamma) + Q(\gamma, R)a \\
\dot{R} &= \bar{M}(\gamma) - R \\
\dot{a} &= a \cdot R^{-1}\frac{\partial \bar{g}(\gamma)}{\partial \gamma} + \frac{1}{2}a'\frac{\partial Q(\gamma, R)}{\partial \gamma}a + \lambda \cdot \bar{M}_\gamma(\gamma) \\
\dot{\lambda} &= \mathcal{H}_R - \lambda \cdot I,
\end{aligned} \tag{5.43}$$

subject to the boundary conditions (5.42), and where $\mathcal{H}_R$ is the derivative of the Hamiltonian with respect to the R matrix.

We have reduced the problem of determining the dominant escape path from a probabilistic problem on a function space of time trajectories to a variational control problem. Further, we have represented the solution of the control problem as a two-point boundary value ODE problem with given initial state conditions and terminal state constraints for an arbitrary time $t < T$. In other words, if we solve (5.43) as a two-point boundary problem with the given initial conditions for γ and R, a terminal condition $\gamma(\tau^e) = \hat{\gamma} \in \partial G$, for some $\tau^e < T$, we will have *an* escape path to $\hat{\gamma}$. We can then evaluate the cost function (5.39) for the arbitrary $(\tau^e, \hat{\gamma})$ to obtain a value (abusing notation slightly) of $S(\tau^e, \hat{\gamma})$. To find the *dominant* escape path, we minimize this function over $\tau^e \in (0, T)$ and over $\hat{\gamma} \in \partial G$:

$$\bar{S} = \inf_{\hat{\gamma} \in \partial G} \inf_{\tau^e \in (0,T)} S(\tau^e, \hat{\gamma}). \tag{5.44}$$

The path that achieves the minimum is the dominant escape path. This path characterizes the evolution of the parameters on the most likely path away from the stable point. The minimized value $\bar{S}$ determines the exponential rate of convergence.

6. NUMERICAL ANALYSIS

6.1. *Numerical results*

In this section we apply the analytic methods from Section 5 to determine the dominant escape paths. We analyse both dynamic and static versions of the model, the distinction being in the specification of X_{n-1} in the government's model. In the static version of the model (2.12) does not contain the lagged variables. For the numerical analysis and simulations, we set the parameters of the Phillips curve to $\theta = 1, u = 5$, and the government's discount factor to $\delta = 0{\cdot}98$. We assumed Gaussian disturbances, and set $(\sigma_1, \sigma_2) = (0{\cdot}3, 0{\cdot}3)$. A simple calculation shows that the self-confirming equilibrium is the intersection of the line $\gamma_1 = \theta$ with the parabola determined by $u - (\gamma_{-1})/(1 + \gamma_1^2) = 0$. There is a unique self-confirming equilibrium, which is depicted in Figure 2. It has $\gamma_{-1} = 10, \gamma_1 = -1$.

As we describe in Appendix E, to compute the dominant escape path numerically, it helps to recast the boundary value problem as an initial value problem. In the ODE system (5.43) and boundaries (5.42), the only components left undetermined are the initial conditions for the co-states. We can solve the problem by minimizing over these initial conditions, and determine the escape times and escape points endogenously. For the escape set G, we take circles of varying radius centred on $\bar{\gamma}$. In Appendix E we briefly describe the numerical techniques used to calculate the dominant escape paths. In order to conserve on dimensionality, in our calculations for the dynamic model we ignore the cross-effects of the

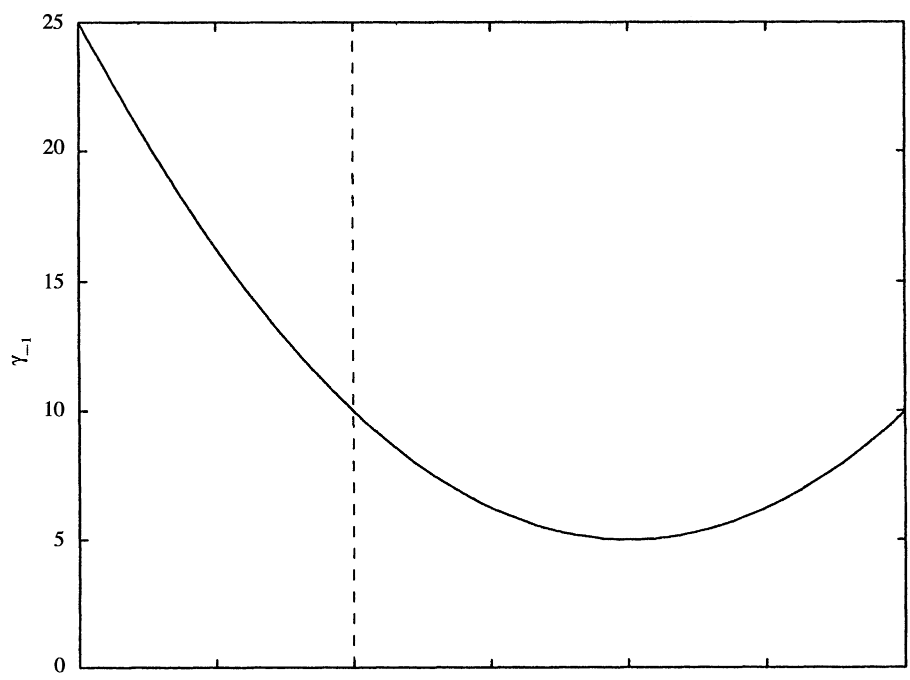

FIGURE 2

Self-confirming equilibrium with $u = 5, \theta = 1$

covariance matrix R. We enforce this by setting the second co-state vector λ in (5.43) identically to zero. The parameter vector γ is six-dimensional, which means that we optimize over the six-dimensional co-state vector α. Additionally, we track the evolution of the upper-triangular elements of R, which is a 21×1 vector. Including the cross-effects would entail optimizing over these additional 21 dimensions which is numerically intractable. In the static model we did allow for the cross-effects, which only slightly changed the dynamics in the time domain.

We turn now to the numerical results. In Figure 3 we report the dominant escape path from the dynamic model, setting the radius of G at 7 Euclidean units. Here we clearly see that along an escape path, the sum of the coefficients on inflation rises from its self-confirming equilibrium value of -1 to (nearly) the induction hypothesis level of zero. By activating the induction hypothesis, the results of Phelps (1967) described above apply and it is optimal for the government to set inflation to (near) zero. Thus, along an escape path the government temporarily learns a version of the theory of the natural rate.

In Figure 4 we plot the escape paths from the static model under the assumptions that the shocks W are distributed normally and binomially, respectively. The escape paths are nearly identical under the two error distributions. This suggests that for the static model it is interesting to consider the simple binomial distribution, which motivates the calculations to be reported in Section 7. In addition, the paths from the static model have many of the qualitative properties of the dynamic model. The escape paths for all of the parameters are characterized by a long period of very slow movement, followed by a rapid change in the

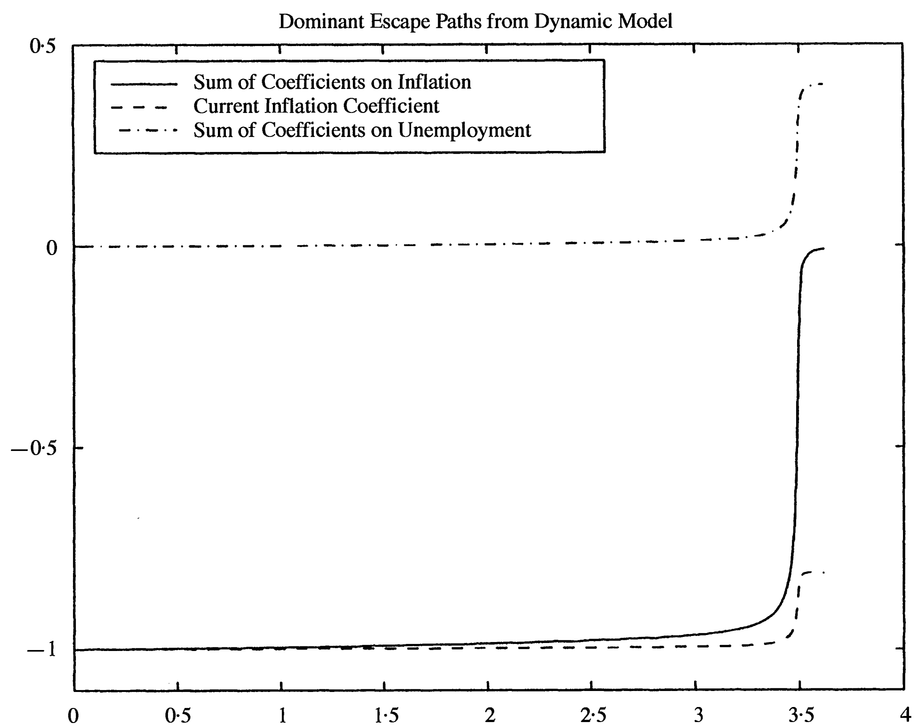

FIGURE 3

Dominant escape path from the dynamic model

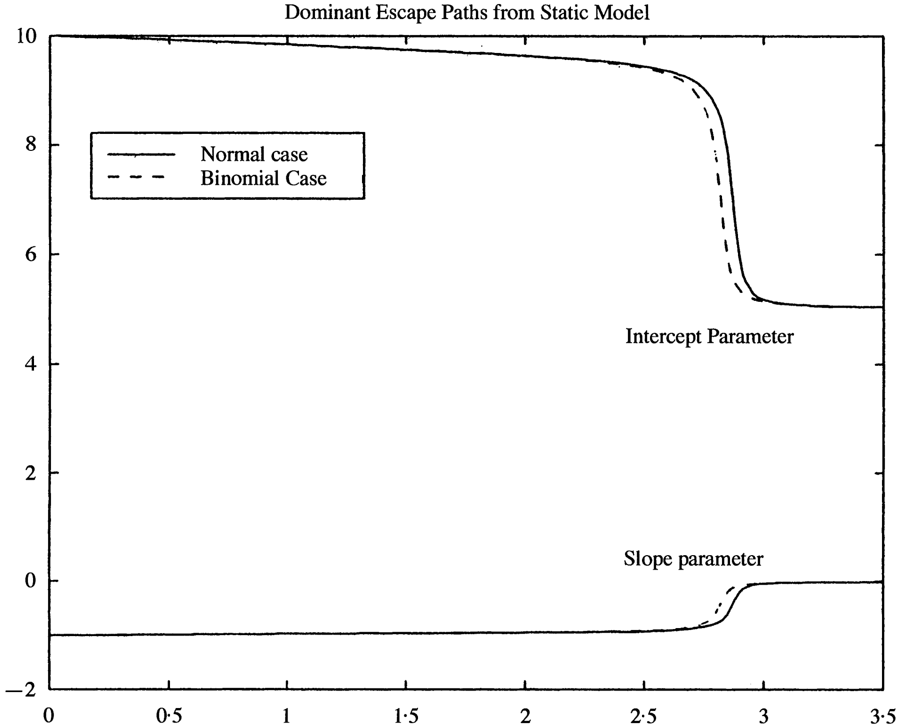

FIGURE 4

Dominant escape path from the static model

same direction. This illustrates the nonlinearity of the differential equations that determine the dominant escape path.

In order to determine the frequency of escapes, we need to consider not only the escape paths, but also the minimized value $\bar{S}$. From Theorem 5.3 above, we know that the value of $\bar{S}$ determines the exponential rate of decay of the probability of an escape. Table 1 presents the rate of convergence $\bar{S}$ and the maximum value along the path of the sum of the coefficients on inflation for different specifications of the model. The table includes estimates of convergence rates for the dynamic model with different size escape sets and for the static model with different error distributions.

TABLE 1

Results for the escape problem

	Radius of set	Sum of coefficients on π	Rate of convergence ($\bar{S}$)
Dynamic model	6·50	−0·0783	$9{\cdot}691\times10^{-6}$
	7·00	−0·0084	$9{\cdot}706\times10^{-6}$
	7·15	−0·0032	$1{\cdot}040\times10^{-5}$
	7·25	−0·0032	$2{\cdot}566\times10^{-4}$
Static model: normal	5·00	−0·0254	$4{\cdot}987\times10^{-4}$
Static model: Binomial	5·00	−0·0257	$4{\cdot}985\times10^{-4}$

A feature of Table 1 that draws immediate notice is that in all specifications the rate of convergence is very slow, which is shown by the low values of $\bar{S}$. This slow convergence rate captures the fact that in our simulations, for any gain setting (no matter how small) we always observed an escape. Also note that, as expected, the rate of convergence increases as we increase the size of the escape set. Clearly the larger the escape set, the less often do we expect escapes. However note that the rate of convergence is nearly the same until the radius is about 7 units, after which it increases dramatically. This reflects the difficulty of pushing beliefs past the Ramsey point, which we discuss more below. Additionally, we see that in the static model not only are the escape paths in the two cases (normal and binomial) nearly identical as in Figure 4, but they also converge at a nearly identical rate.

6.2. *Mean escape times in static and dynamic models*

In the figures and the table, we also note some important differences between the dynamic and the static model. Although the two specifications share the same self-confirming equilibrium, they differ out of equilibrium along an escape path. In the table we see that the rates of convergence differ by an order of magnitude between the dynamic model and the static model. This suggests that by allowing more flexibility in the specification of the government's beliefs, we enable policymakers more rapidly to detect the induction hypothesis. The importance of the lags in the escape paths can also be seen by comparing the paths for the sum of the coefficients on inflation with the coefficient on current inflation. In Figure 4, we see that the escape paths in the static model cause the government's perceived Phillips curve to become vertical. This specification trivially satisfies the induction property, and it would be possible that the same dynamics would recur in the full dynamic model. However in Figure 3 we see that in the dynamic model, along an escape most of the movement in the sum of the coefficients on inflation is due to the coefficients on the lags. Thus in the dynamic model, even on an escape there is still a short-run Phillips curve, but the government comes to believe that the long-run Phillips curve is vertical, and thus it sets inflation to zero.

6.3. *Escape time distributions*

Next we compare our predictions on the escape problem to some additional results from simulations, and we see that our calculations provide an accurate description of the escape problem. Theorem 5.3 above suggests that the asymptotic distribution of escape times is exponential, and this appears to be roughly borne out by the empirical distribution that we plot in the top panel of Figure 5. The figure shows a histogram of the distribution of the time of first escape from the self-confirming equilibrium in the dynamic model for 1000 simulated paths, with the gain set at the very low value $\varepsilon = 0{\cdot}001$. Here we see that the distribution is clearly skewed and has a long tail, resembling an exponential distribution. We have also seen that the mean escape times increase (at least) exponentially as the gain decreases, and this is shown in the bottom panel of Figure 5. The figure plots the mean time of first escape from the self confirming equilibrium for 1000 simulated paths in both the static model and the dynamic model for different gain settings.[12] We have also seen in Table 1 that the rate of convergence is much slower in the dynamic model, and this is also

12. We set the radii of G so that the boundary of G is close to the Ramsey point. Thus we use a radius of 5 units in the static model and 7 units in the dynamic model.

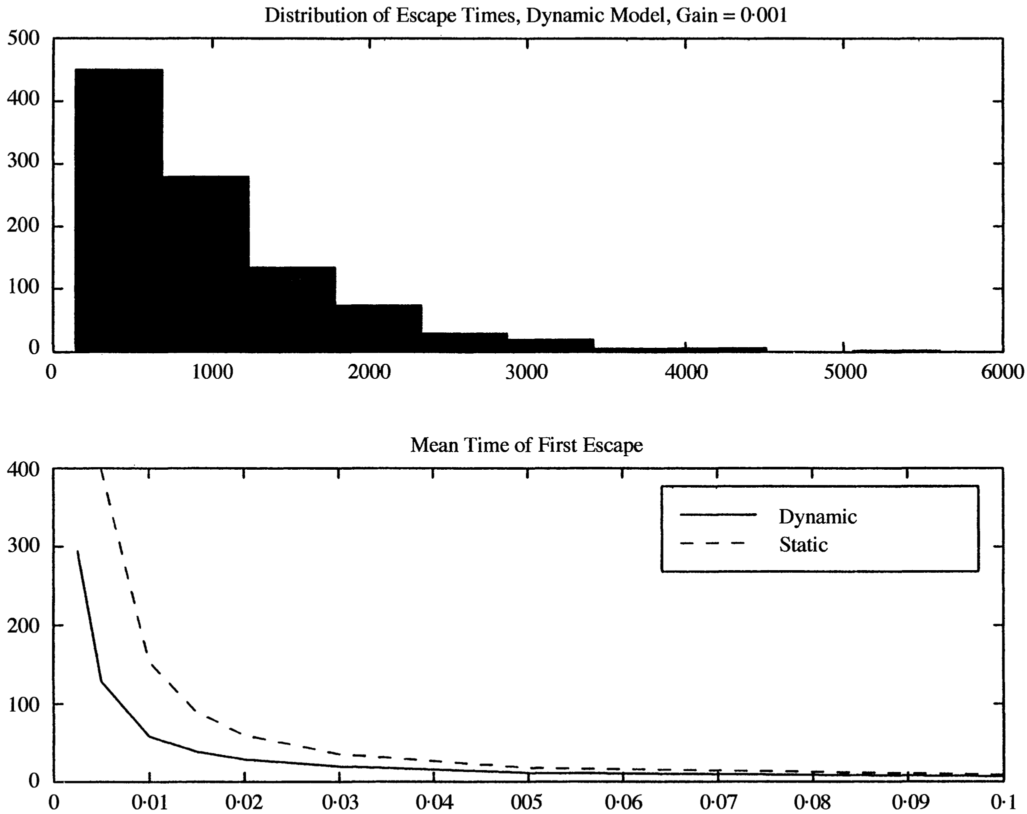

FIGURE 5

Simulation results from the dynamic and static models

clear from the figure. As the gain decreases, the escapes become exponentially more rare, but occur much more frequently in the dynamic model.

The dominant escape paths that we have derived and calculated in this paper also describe the escape dynamics from simulations like those of Sims (1988), Chung (1990), Sargent (1999). In Figure 6, we see that our calculations match very closely the direction in which the government's beliefs escape the self-confirming equilibrium. In the figure, we project the escape paths onto a subset of the parameter space and plot the sum of the coefficients on inflation vs. the constant coefficient. The dominant escape path is shown by the nearly straight line that leads from the self-confirming values $(10, -1)$ in the upper left of the figure downward to the lower right. Also shown is the simulated path from Figure 5 that has the escape time closest to the sample mean. The figure clearly shows that the mean simulated path lies within a small neighbourhood of the dominant path. Indeed *all* of the simulated paths are within a small neighbourhood of the dominant path.

In addition to our quite accurate results in the parameter space, our calculations provide a good description of the escape dynamics in the time domain. In Figure 7 we plot the dominant path escape paths for the static model, with the time axis scaled to $\log(t/\varepsilon)$. We also plot some results from the 1000 simulated escape paths with gain $\varepsilon = 0{\cdot}0025$. Here we show the escape paths with the minimum, maximum, and mean

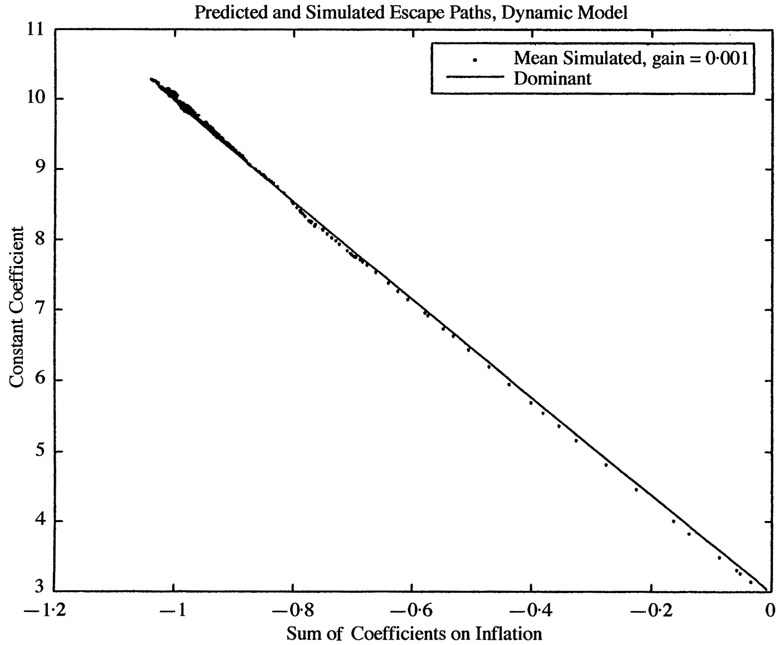

FIGURE 6

Dominant and simulated escape paths from the dynamic model

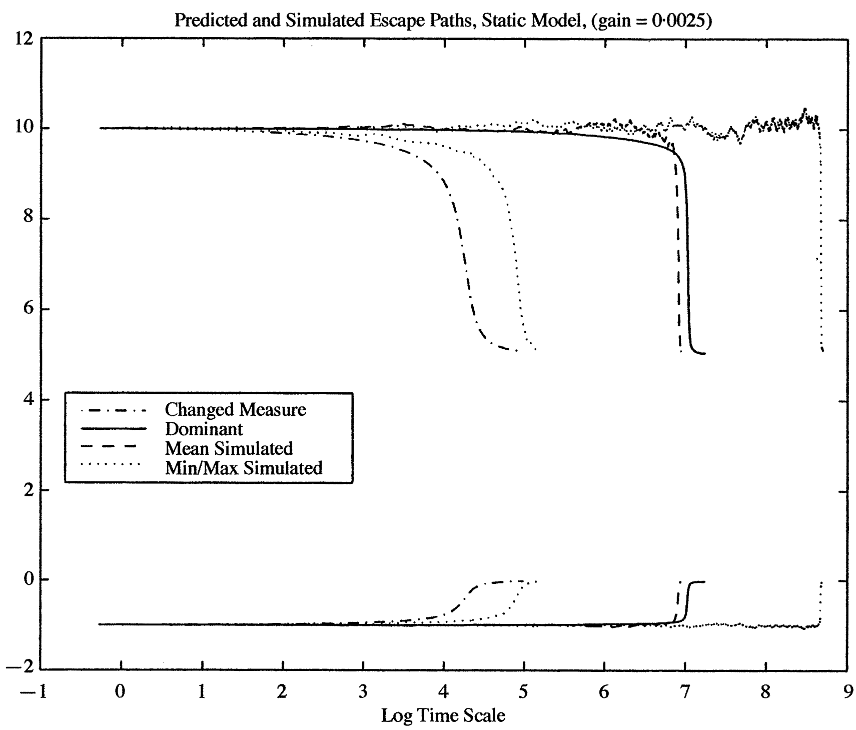

FIGURE 7

Dominant and simulated escape paths from the static model escape times. Also for reference, we plot the escape path that results from the analysis of Section 7, in which we transform probability measures. The figure clearly shows that the dominant path lies within the band of simulated outcomes, and that the dominant escape path is quite close to the mean observed path. As we shall discuss, the analysis of Section 7 leads to the same direction of escape, but the time dynamics are distorted by the change in measures. In the figure we see that this escape path is shorter than even the shortest of the 1000 simulated paths. As the gain decreases, the band of outcomes narrows further and the mean path gets closer and closer to the dominant path. Thus the numerical implementations of the analytic characterizations provide very accurate estimates of the convergence and escape properties of monetary policymakers' beliefs.

6.4. *Directions of mean dynamics along the escape path*

Because the mean dynamics are locally stable around the unique self-confirming equilibrium, the escape dynamics are essential in starting departures from the self-confirming equilibrium $\bar{\gamma}$. However, for our model the mean dynamics themselves have features that promote rapid movements toward the Ramsey outcome after the escape dynamics have initiated an escape from a neighbourhood of $\bar{\gamma}$. Thus, while the mean dynamics are both locally and globally stable around the self-confirming equilibrium, beyond a particular neighbourhood of a self-confirming equilibrium, they point *toward* the beliefs supporting the Ramsey outcome. For the static ($X_{n-1} = 1$) model, we illustrate this feature in Figures 8 and 9, which show both the mean dynamics and the escape dynamics $\dot{v}$ along the dominant escape path. In Figure 8 we plot the dynamics over the entire dominant escape

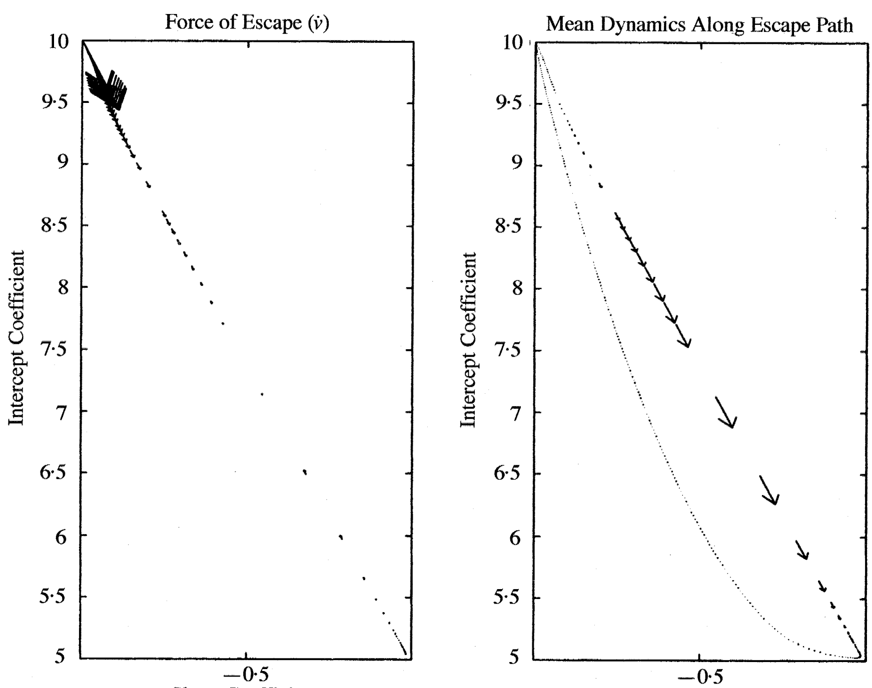

FIGURE 8

The force of escape and mean dynamics along the dominant escape path, static model

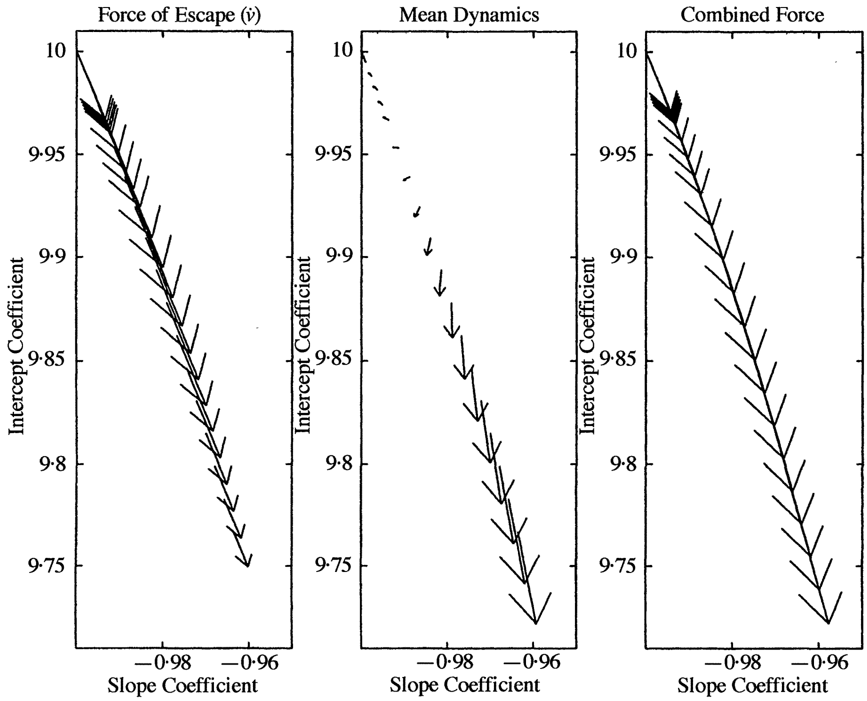

FIGURE 9

Detail of the force of escape, mean dynamics, and combined force in a neighbourhood of the self-confirming equilibrium

path, and Figure 9 shows a closeup of the escape, mean, and total dynamics in a neighbourhood of the self-confirming equilibrium. As Figure 9 shows, near the self-confirming equilibrium, the mean dynamics point toward the self-confirming equilibrium and the escape dynamics $\dot{v}$ point away (toward beliefs that support Ramsey). In this neighbourhood, the mean dynamics and escape dynamics oppose each other. The mean dynamics prevail under the natural distribution of shocks. However, outside of this neighbourhood, the mean dynamics reinforce the escape dynamics by pointing toward Ramsey, as Figure 8 further illustrates. This figure shows that the magnitude of $\dot{v}$ essentially falls to zero after the first instants, and the mean dynamics start to push toward Ramsey with greater force. Thus, if we were to initiate government beliefs along the escape path sufficiently far from the self-confirming equilibrium, the mean dynamics themselves would sweep beliefs toward the beliefs that support the Ramsey outcome. The right panel of Figure 8 also shows the mean dynamics *from* the Ramsey outcome to the self-confirming equilibrium. They complete the circuitous path of the mean dynamics which (starting along the escape path) send the system to the Nash outcome by travelling near the Ramsey outcome.[13]

13. Evans and Honkapohja (2001) had noted such behaviour of the mean dynamics for this model.

6.5. *Escaping the experimentation trap*

Within the confines of the government's approximating model, detecting the natural rate hypothesis requires that there be sufficient dispersion in the public's expected rate of inflation. But within a self-confirming equilibrium, there is *no* variation in the expected rate of inflation because the government does not vary its setting of the systematic part of inflation x_t. Though the outcome is the same, the structure of this experimentation trap differs from the trap in Kydland and Prescott's time consistent equilibrium. Here the government fails to generate the range of experiments needed to detect the natural rate hypothesis within its approximating model. But only if it detects something approximating the natural rate hypothesis will it want to generate those experiments.

Along the escape route the government generates those experiments. The experiments are initiated by an unusual shock process encoded in $\dot{v}$, to be analysed in more detail in Section 7. In the static model, any force that causes the government to experiment by randomizing x_t generates a (y_t, U_t) data scatter through (2.8) that *steepens* (in the (y, U) plane) the estimated Phillips curve. Through the government's best response map, any steepening of the Phillips curve causes the government to lower inflation, generating influential observations that make the Phillips curve steeper. Overweighting recent observations helps this process along. This self-reinforcing process comes to a halt when the estimated Phillips curve becomes vertical. The system cannot remain at the Ramsey outcome forever, because there is in truth a short-run Phillips curve that the government will discover and begin to exploit, rekindling the mean dynamics that drive the system toward the Nash outcome.

7. INTERPRETING $\dot{v}$ VIA ORTHOGONALITY CONDITIONS

By studying the special case of the static model ($X_{n-1} = 1$) and binomial shocks, this section identifies the escape route using a different argument than Section 5. The argument works directly with the least squares orthogonality conditions that drive the government's beliefs, and finds the appropriate conditional distribution of the shocks that gives rise to the escape dynamics.

7.1. *Analysing the orthogonality conditions*

The mean dynamics (4.32) come from writing the recursive least squares orthogonality condition as

$$\frac{\gamma_{n+1} - \gamma_n}{\varepsilon} = R_n^{-1}\bar{g}(\gamma_n) + R_n^{-1}(g(\gamma_n, \xi_n) - \bar{g}(\gamma_n)), \tag{7.45}$$

then using a limiting argument and a law of large numbers to replace the term $(g(\gamma_n, \xi_n) - \bar{g}(\gamma_n))$ with its unconditional mean of 0 while driving ε to zero. The escape dynamics also originate from (7.45). However, problem (5.39) induces a different average "forcing function" than the identically zero $R_n^{-1}(Eg(\gamma_n, \xi_n) - \bar{g}(\gamma_n))$. Instead, problem (5.39) averages not with respect to the unconditional distribution of ξ_n but with respect to a distribution whose conditional expectation we denote $\tilde{E}$. This makes the limiting escape dynamics come from

$$\frac{\gamma_{n+1} - \gamma_n}{\varepsilon} \approx R_n^{-1}\bar{g}(\gamma_n) + R_n^{-1}\tilde{E}(g(\gamma_n, \xi_n) - \bar{g}(\gamma_n)) \tag{7.46}$$

or

$$\dot{\gamma} = R^{-1}\bar{g}(\gamma) + \dot{v} \tag{7.47}$$

where

$$\dot{v} \approx R_n^{-1}\tilde{E}(g(\gamma_n, \xi_n) - \bar{g}(\gamma_n)). \tag{7.48}$$

We want to shed light on the probability distribution of shocks inducing $\tilde{E}$. For a W that is jointly binomial, we can find a transformed measure that gives us a very convenient interpretation of $\dot{v}$ in (7.47). Our method is to work directly with the time n orthogonality conditions (4.24), (4.25) that dictate the movement of γ_n.

Throughout this section, we assume that W_{in} has the binomial distribution with variance 1 for $i = 1, 2$:

$$W_{in} = \begin{cases} 1 & \text{with probability } \frac{1}{2} \\ -1 & \text{with probability } \frac{1}{2}. \end{cases} \tag{7.49}$$

For simplicity, we examine the static model in which

$$X_{n-1} = 1 \quad \forall n \geqq 1$$

so that the best response mapping of the government is

$$h(\gamma_1, \gamma_{-1}) = -\frac{\gamma_1 \gamma_{-1}}{1 + \gamma_1^2}.$$

Under this specification, we shall find a transformed measure that describes the escape point[14] and according to which

$$\dot{v} = R^{-1}\begin{bmatrix} \sigma_1\sigma_2 \\ 0 \end{bmatrix}.$$

We obtain this representation by finding a most likely unlikely shock sequence that moves the government's beliefs a given distance away from a self-confirming equilibrium. Our method of analysis is basically to study the different possible unlikely sequences of shocks that move the government's beliefs a given distance away from the self-confirming equilibrium $\bar{\gamma}$, and among these to find the most likely sequence. The static model with the binomial shocks[15] is simple enough to let us get our hands on these sequences of shocks. The heart of the argument is to notice that the most likely of such unlikely sequences has the property that each realization of the shocks must push the learning algorithm in the same direction away from $\bar{\gamma}$.

The mean dynamics and the escape dynamics both originate from the same stochastic difference equation system (4.24), (4.25). Above, we have defined η_n as the residual in the government's Phillips curve for regression coefficients γ and shock vector W_n. Recall that

$$g(\gamma_n, \xi_n) = \eta_n \begin{bmatrix} \pi_n \\ X_{n-1} \end{bmatrix},$$

14. The transformed measure underestimates the time to escape, as indicated in Section 6.
15. And more generally, the static model with multinomial shocks analysed by Cho and Sargent (1999).

where η_n is the forecasting error. To emphasize the link between the dynamics of γ_n and the original perturbation W_n, we write

$$\tilde{g}(\gamma_n, \sigma W_n) = g(\gamma_n, \xi_n),$$

and

$$\eta_n = \eta(\gamma_n, \sigma W_n)$$

where $\sigma W_n = (\sigma_1 W_{1,n}, \sigma_2 W_{2,n})$. Let $w \in \{-1, 1\}^2$ be a particular realization of W_n. Since W_n can take four different values, for any $\gamma \in \Re^2$, $\{\tilde{g}(\gamma, \sigma W_n)\}$ consists of four vectors. Let $R_n = [R_{ij,n}]$ and D_n be the determinant of R_n. We can write

$$\begin{bmatrix} \gamma_{1,n+1} \\ \gamma_{-1,n+1} \end{bmatrix} = \begin{bmatrix} \gamma_{1,n} \\ \gamma_{-1,n} \end{bmatrix} + \frac{\varepsilon}{D_n} \begin{bmatrix} -R_{21,n} + (x_n + \sigma_2 W_{2,n}) \\ R_{11,n} - R_{21,n}(x_n + \sigma_2 W_{2,n}) \end{bmatrix} \eta_n. \tag{7.50}$$

For each realization $w \in \{-1, 1\}^2$ of W_n it is useful to depict the contour of γ satisfying $\eta(\gamma, \sigma W_n) = 0$:

$$\{\gamma \in \Re^2 : \eta(\gamma, \sigma w) = 0\}.$$

A self-confirming equilibrium is a point that in effect is the unconditional average over these contours. As depicted in Figure 10, $\gamma^{1,e}$ is the intersection of

$$\{\gamma : \eta(\gamma, (\sigma_1, \sigma_2)) = 0\}$$

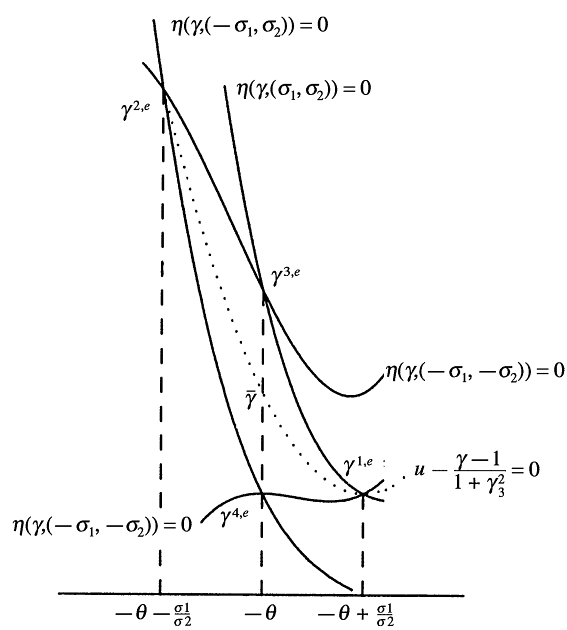

FIGURE 10

The horizontal axis is γ_1, and the vertical axis is γ_{-1}

and

$$\{\gamma : \eta(\gamma, (-\sigma_1, -\sigma_2)) = 0\}.$$

Note that the stable point $\bar{\gamma}$ is surrounded by the four curves.

For convenience, we choose the area surrounded by the four curves (including its boundary) as G. Note that ∂G is formed by the union of the segments of the four curves conditioned on four different realizations of W_n. The characterization of the most likely exit point with respect to a general compact convex neighbourhood of $\bar{\gamma}$ follows from the same logic.

Since $\tilde{E}$ is induced by the probability distribution over the set of escape paths, we need to examine the probability distribution of the exit points along the boundary of the neighbourhood of $\bar{\gamma}$. However, the entire set of escape paths from the stable point to a boundary of G is too complicated to be characterized directly, because there are numerous ways to escape from the stable point to a fixed point in ∂G as indicated (5.44). Thus, we proceed by constructing a subset of escape paths that are analytically manageable, while inducing a probability measure that is absolutely continuous with respect to the original probability measure over the set of all exit points. Theorem 5.3 indicates that the probability distribution over exit points collapses to the exit point from the dominant escape path as the gain ε converges to 0. Thus, if we can construct an absolutely continuous probability measure over the set of exit points, then the new measure must converge to the same exit point.

To construct a transformed measure with the desired property, we begin by noting that from (7.50)

$$\frac{d\gamma_{-1,n+1}}{d\gamma_{1,n+1}} = \frac{R_{11,n} - R_{21,n}(x_n + \sigma_2 W_{2,n})}{-R_{21,n} + (x_n + \sigma_2 W_{2,n})} \tag{7.51}$$

is independent of $W_{1,n}$, which implies that for any W_n, $\{R_n^{-1}\tilde{g}(\gamma_n, \sigma W_n)\}$ consists of two pairs of linearly dependent vectors, each pair being indexed by one of the two possible values for $\sigma_2 W_{2n}$. Thus in (7.50), the evolution of $\gamma_{1,n}$ and $\gamma_{-1,n}$ is influenced by a common scalar factor η_n. By (7.51), the vector field $\{R_n^{-1}\tilde{g}(\bar{\gamma}, \sigma W_n)\}$ around the stable point $\bar{\gamma}$ consists of two *pairs* of linearly dependent vectors. In particular, (7.51) implies that

$$\exists\rho \neq 0, \quad \text{such that} \quad R_n^{-1}\tilde{g}(\gamma, (\sigma_1, \sigma_2)) = \rho R_n^{-1}\tilde{g}(\gamma, (-\sigma_1, \sigma_2)) \tag{7.52}$$

$$\exists\rho' \neq 0, \quad \text{such that} \quad R_n^{-1}\tilde{g}(\gamma, (\sigma_1, -\sigma_2)) = \rho' R_n^{-1}\tilde{g}(\gamma, (-\sigma_1, -\sigma_2)). \tag{7.53}$$

Under the "usual" event, each element $W_n \in \{-1, 1\}^2$ is realized with probability 1/4. Because $\bar{\gamma}$ is the stable point of the (unconditional) mean dynamics,

$$\sum_{w\in\{-1,1\}^2} \tfrac{1}{4} R_n^{-1}\tilde{g}(\bar{\gamma}, \sigma w) = 0.$$

Therefore, the two pairs of linearly dependent vectors must point to opposite directions:

$$\rho < 0 \quad \text{and} \quad \rho' < 0. \tag{7.54}$$

Otherwise, the mean dynamics cannot be zero as required at $\bar{\gamma}$.

Since we can uniquely identify the sequence of perturbations for each escape path, it makes sense to say that an escape path is generated by the sequence of perturbations.

Definition 7.1. *We say that $w, w' \in \{-1, 1\}^2$ satisfies the same half-space condition if $R_n^{-1}\tilde{g}(\gamma_n, \sigma w)$ and $R_n^{-1}\tilde{g}(\gamma_n, \sigma w')$ are located in the same open half space. We say that an escape path with exit time τ^e satisfies the same half-space condition, if it is generated by perturbations that satisfy the same half-space condition for all $\tau \in (0, \tau^e)$.*

The same half-space condition requires that the two vectors $R_n^{-1}\tilde{g}(\gamma_n, \sigma w)$ and $R_n^{-1}\tilde{g}(\gamma_n, \sigma w')$ are pointing in the "same" direction. In the case of binomial distributions, w and w' satisfy this condition unless $R_n^{-1}\tilde{g}(\gamma_n, \sigma w)$ and $R_n^{-1}\tilde{g}(\gamma_n, \sigma w')$ are linearly dependent vectors that point in opposite directions.

The law of large numbers says that under the "usual" events, each element in $w \in \{-1, 1\}^2$ occurs with 1/4 of frequencies over $(0, \tau)$ for any $\tau > 0$. If so, $\dot{v}$ must be 0 over the same interval of time and γ moves toward $\bar{\gamma}$. Thus, any move away from the stable point $\bar{\gamma}$ requires a sequence of "unusual" events of W_n, whose probability is strictly less than 1. In order to maximize the probability of escape through a *particular* point at the boundary of G, the escape path must minimize the number of "unusual" events to reach the neighbourhood of the exit point. In order to reach a particular point on the boundary in the most economical way, the adjustment term $R_n^{-1}\tilde{g}(\gamma_n, \sigma w)$ associated with each realization w of W_n must point to the same "direction". Recall that $\{R_n^{-1}\tilde{g}(\gamma_n, \sigma w)\}$ are two pairs of linearly dependent vectors and that the two vectors in each pair point in opposite directions. If the same half-space condition is violated, then some moves cancel others, wasting precious time to escape. Therefore, using the same logic as in (5.44) above, if $\gamma(\tau)$ is an escape path through a particular point $\hat{\gamma} \in \partial G$ with exit time τ^e:

$$\gamma(0) = \bar{\gamma} \quad \text{and} \quad \gamma(\tau^e) = \hat{\gamma} \in \partial G$$

that takes the minimum amount of time among all escape paths through $\hat{\gamma}$, then the escape path should be generated by sequence of perturbations satisfying the same half-space condition for all $\tau \in (0, \tau^e)$.

Lemma 7.2. *Fix $\hat{\gamma} \in \partial G$, and let γ be the escape path through $\hat{\gamma}$ that minimizes the exit time among all escape paths through $\hat{\gamma}$.*

1. *γ must be generated by a sequence of perturbations that satisfy the same half-space condition.*
2. *If an escape path γ satisfies the same half-space condition, then no more than two perturbations can be realized along γ.*

Proof. See Appendix D. ||

Let $\mathbf{D}^\varepsilon$ be the set of all escape paths when the gain is $\varepsilon > 0$. Instead of all escape paths, let us consider the set $\mathbf{D}_s^\varepsilon$ of all escape paths that satisfy the same half-space condition. Because W_n can take four different values, we can make six different pairs of the realized values of W_n. Among the six pairs, $\{(1, 1), (-1, 1)\}$ and $\{(1, -1), (-1, -1)\}$ induce a pair of $R_n^{-1}\tilde{g}(\gamma_n, \cdot)$ that are linearly dependent. The remaining four pairs, namely (1) $\{(1, 1), (-1, -1)\}$, (2) $\{(1, -1), (-1, 1)\}$, (3) $\{(1, 1), (1, -1)\}$ and (4) $\{(-1, 1), (-1, -1)\}$, satisfy the same half-space condition.

The next result is crucial.

Proposition 7.3. *Fix a measurable $A \subset \partial G$. If $\lim_{\varepsilon \to 0} \Pr(A : \mathbf{D}^\varepsilon) = 0$, then*

$$\lim_{\varepsilon \to 0} \Pr(A : \mathbf{D}_s^\varepsilon) = 0.$$

Proof. See Appendix D. ||

Because we know that Pr $(\cdot : \mathbf{D}^{\varepsilon})$ is concentrated on the exit point of the dominant escape path, Proposition 7.3 allows us to find the most likely exit point by using $\mathbf{D}_s^{\varepsilon}$ instead of $\mathbf{D}^{\varepsilon}$.

Recall that $\mathbf{D}_s^{\varepsilon}$ admits an escape path that is generated by $\{(1, 1), (-1, -1)\}$, for example, in which $(1, 1)$ is realized with frequency p and $(-1, -1)$ with frequency $1 - p$ for $p \in [0, 1]$. As $\varepsilon \to 0$, the law of large numbers implies that the most likely frequency of $(1, 1)$, conditioned on that path is generated by $\{(1, 1), (-1, -1)\}$, is

$$\Pr(W_n = (1, 1) : W_n \in \{(1, 1), (-1, -1)\}) = \tfrac{1}{2}.$$

Thus, if the probability distribution over $\mathbf{D}_s^{\varepsilon}$ is collapsed to a single point, then the point should be the exit point of the escape path generated by one of the above four combinations of perturbations, in which each element is realized with probability one half. By using the standard result of the stochastic approximation, we can represent these four escape paths in terms of an ODE calculated with respect to the conditional probability distribution of W_n.

7.2. *A "race" of four ODEs*

We construct four ODEs associated with the four conditional distributions satisfying the same half-space condition. The one that heads toward the boundary most quickly corresponds to the dominant escape path. First, if $W_n \in \{(1, 1), (-1, -1)\}$, then the associated conditional ordinary differential equation is

$$\dot{\gamma} = R^{-1} \begin{bmatrix} -\dfrac{\gamma_{-1}\gamma_1}{1+\gamma_1^2}\left(u - \dfrac{\gamma_{-1}}{1+\gamma_1^2}\right) - (\theta + \gamma_1)\sigma_2^2 + \sigma_1\sigma_2 \\ u - \dfrac{\gamma_{-1}}{1+\gamma_1^2} \end{bmatrix} \tag{7.55}$$

which has the stable point

$$\gamma^{1,e} = \left(-\theta + \frac{\sigma_1}{\sigma_2}, u\left[-\theta + \frac{\sigma_1}{\sigma_2}\right]^2\right),$$

depicted in Figure 10. It is important to note that the R matrix also evolves along the candidate escape paths. It's differential equation is in general given by the conditional expectation of M_n as in (3.19), conditioned on the half space shock realizations. However, in the static case, M_n is independent of W_{1n}, and along the four candidate paths the marginal distribution of W_{2n} is equal to its unconditional distribution. Therefore we have that, as in the analysis in Section 5, along the escape paths R follows the mean dynamics (4.33) conditional on γ. Because our focus is on the escape properties of γ, we shall suppress the evolution of the covariance matrix.

If $W_n \in \{(1,-1),(-1,1)\}$, then the associated conditional ordinary differential equation is

$$\dot{\gamma} = R^{-1}\begin{bmatrix} -\frac{\gamma_{-1}\gamma_1}{1+\gamma_1^2}\left(u - \frac{\gamma_{-1}}{1+\gamma_1^2}\right) - (\theta+\gamma_1)\sigma_2^2 - \sigma_1\sigma_2 \\ u - \frac{\gamma_{-1}}{1+\gamma_1^2} \end{bmatrix} \tag{7.56}$$

which has the stable point

$$\gamma^{2,e} = R^{-1}\left(u\left(1+\left[-\theta-\frac{\sigma_1}{\sigma_2}\right]^2\right), -\theta-\frac{\sigma_1}{\sigma_2}\right).$$

If $W_n \in \{(1,1),(1,-1)\}$, then the associated conditional ordinary differential equation is

$$\dot{\gamma} = R^{-1}\begin{bmatrix} -\frac{\gamma_{-1}\gamma_1}{1+\gamma_1^2}\left(u - \frac{\gamma_{-1}}{1+\gamma_1^2}+\sigma_1\right) - (\theta+\gamma_1)\sigma_2^2 \\ u - \frac{\gamma_{-1}}{1+\gamma_1^2}+\sigma_1 \end{bmatrix} \tag{7.57}$$

which has the stable point

$$\gamma^{3,e} = ((u+\sigma_1)(1+\theta^2), -\theta).$$

If $W_n \in \{(-1,1),(-1,-1)\}$, then the associated conditional ordinary differential equation is

$$\dot{\gamma} = R^{-1}\begin{bmatrix} -\frac{\gamma_{-1}\gamma_1}{1+\gamma_1^2}\left(u - \frac{\gamma_{-1}}{1+\gamma_1^2}-\sigma_1\right) - (\theta+\gamma_1)\sigma_2^2 \\ u - \frac{\gamma_{-1}}{1+\gamma_1^2}-\sigma_1 \end{bmatrix} \tag{7.58}$$

which has the stable point

$$\gamma^{4,e} = ((u-\sigma_1)(1+\theta^2), -\theta).$$

Notice the relationship between $\gamma^{j,e}$ $(j=1,\ldots,4)$ and the perturbations that generate each ordinary differential equation. For example, (7.55) is generated by $W_n \in \{(1,1),(-1,-1)\}$ which has $\gamma^{1,e}$ as its stable point, where $\gamma^{1,e}$ is the intersection of

$$\{\gamma : \eta(\gamma,(\sigma_1,\sigma_2)) = 0\}$$

and

$$\{\gamma : \eta(\gamma,(-\sigma_1,-\sigma_2)) = 0\}.$$

Our remaining task is to identify the most likely exit point among $\{\gamma^{j,e} : j=1,\ldots,4\}$, each of which is located on the boundary of G that is surrounded by the four curves depicted in Figure 10. Notice that each curve is generated by exactly two perturbations, each of which has the same probability 1/4 of realization. Thus, if it takes $T(\varepsilon)$ periods for the escape path to reach the boundary of G, then the probability that the exits occurs through the point is proportional to $4^{-T(\varepsilon)}$. Moreover, $T(\varepsilon)$ increases as $\varepsilon \to 0$. Thus, the exit point that can be reached in the shortest periods is the most likely exit point among $\{\gamma^{j,e} : j = 1, \ldots, 4\}$. A simple numerical analysis reveals that $\gamma^{1,e}$ is the most likely exit point, because the velocity vector of (7.55) is substantially larger than that of the other three ODEs. Then, by invoking Proposition 7.3, we conclude that this is the point through which the dominant escape exits G. See Cho and Sargent (1999) for additional details.

In conclusion, under the binomial assumption on W_n, the "unconditional" ODE for γ is

$$\bar{g}(\gamma) = R^{-1} \begin{bmatrix} -\dfrac{\gamma_{-1}\gamma_1}{1+\gamma_1^2}\left(u - \dfrac{\gamma_{-1}}{1+\gamma_1^2}\right) - (\theta + \gamma_1)\sigma_2^2 \\ u - \dfrac{\gamma_{-1}}{1+\gamma_1^2} \end{bmatrix},$$

which implies that (7.55) can be written as

$$\dot{\gamma} = \bar{g}(\gamma) + R^{-1} \begin{bmatrix} \sigma_1\sigma_2 \\ 0 \end{bmatrix}$$

and

$$\dot{v} = R^{-1} \begin{bmatrix} \sigma_1\sigma_2 \\ 0 \end{bmatrix}.$$

8. CONCLUDING REMARKS

Fudenberg and Levine (1993), Fudenberg and Kreps (1995), and Fudenberg and Levine (1998) have recommended the self-confirming equilibrium concept partly because of its status as the limit point of a class of what Fudenberg and Kreps (1995) call "asymptotically myopic" learning schemes. We believe that self-confirming equilibria are useful tools for macroeconomics, where there have always been controversies about whether the government's model is specified properly, and where there is a long tradition of academics trying to improve the government's model specification. The Phillips curve example of Kydland and Prescott (1977) is just one example in this tradition.

Macroeconomic examples of self-confirming equilibria typically have the structure that the beliefs of a large player (namely, the government) influence stochastic processes of outcomes significantly. Where the beliefs of large players matter, adaptive learning schemes, like those analysed here and in Fudenberg and Kreps (1995) and Sargent (1999), let escape dynamics have big and recurrent effects on outcomes. In addition to developing some of the analytic results applied in this paper, Williams (2001) provides examples of escape dynamics in adaptive models of oligopoly and growth with production externalities. Also see Bullard and Cho (2001).

In this paper, we have shown that the escape dynamics exhibit the same "near determinism" (in the sense of Whittle (1996)) as the mean dynamics already familiar from the literature on least squares learning. Our numerical calculations of escape probabilities closely fit numerical simulations of our model. An interesting feature of the calculations is how the escape probabilities are larger when the government's model is more richly specified, permitting it to discover the subtler "induction-hypothesis" version of the natural rate hypothesis, despite the fact that the generality of this model relative to the "static" model adds nothing *within* the self-confirming equilibrium.

APPENDIX A. PROOF OF THEOREM 4.1

The result follows directly from Theorem 8.5.1 in Kushner and Yin (1997), under Assumptions A above. The theorem requires the additional assumptions (A8.1.9), (A8.5.2), (A8.5.3) and (A8.5.5) which hold trivially here, since $E_n Z_n = \bar{b}(\theta_n^\varepsilon)$. This implies that the limit in (A8.1.9) is identically zero, that the β_n^ε terms in (A8.5.2) and (A8.5.5) are also identically zero, and that (A8.5.3) is equivalent to (A8.5.1). The theorem is also stated under the weaker condition (A8.5.4), which is implied by (A8.5.4a) above.

We now verify that Assumptions A hold when the shocks W are i.i.d. normal. For the normal case, (A8.1.6) clearly holds by inspection. It is clear that there is a unique locally stable point of the ODE, and furthermore it can be shown that the domain of attraction of the ODE is the entire space, so (A8.5.4a) holds. For (A8.5.0), notice that by the independence of the shocks, we can write

$$P(|\theta_n^\varepsilon| \geqq K) = P(f(w) \geqq K)P(g(w) \geqq K),$$

where w has a standard normal distribution and f and g are some quadratic functions. Denoting the roots of $f(w) - K$ and $g(w) - K$ as (f_1, f_2) and (g_1, g_2) respectively, we have

$$P(|\theta_n^\varepsilon| \geqq K) = (\Phi(f_1) + 1 - \Phi(f_2))(\Phi(g_1) + 1 - \Phi(g_2)),$$

where Φ is the standard normal c.d.f. Tightness follows by noting that the absolute values of the roots are increasing in K and taking limits as $K \to \infty$. For (A8.5.1) and (A8.5.3), note that $|Z_n^\varepsilon|^2$consists of normally distributed random variables up to the fourth order, and so have finite expectation, which implies the uniform integrability. Finally (A8.1.7) holds because υ_n^ε consists of normally distributed random variables up to the second order, and thus can be bounded to arbitrary accuracy on an appropriate compact set. ||

APPENDIX B: PROOF OF THEOREM 5.3

B.1. Part 1. The result follows from Dupuis and Kushner (1989), Theorem 3.2, which requires that paper's assumptions 2.1–2.3 and 3.1. Their assumption 2.2 is satisfied by Assumptions A, and 2.3 is not necessary in the constant gain case, as we restrict our analysis to a finite time interval. Assumption 3.1 is satisfied by our definition of S above. All that remains is 2.1. Under the exponential tail condition given in Part 1, Dupuis and Kushner (1989) Theorem 7.1 (with special attention to the remarks following it) and their Example 7.1 show that 2.1 holds.

B.2. Part 2. The result is a simple application of Kushner and Yin (1997) Theorem 6.10.1, whose assumptions follow directly under the boundedness assumption. The identification of the H function follows from Dupuis and Kushner (1989) Theorems 4.1 and 5.3.

B.3. Parts 3 and 4. After establishing part 2, these results follow directly from Dembo and Zeitouni (1998) Theorem 5.7.11. ||

APPENDIX C: CONSTRUCTING THE Q MATRIX

Recall that in the full dynamic model, the government's regression equation includes current inflation, two lags of inflation and unemployment, and a constant. Therefore γ is a (6×1) vector, and so Q is a (6×6) matrix. The matrix has the form:

$$Q(\gamma, R) = R^{-1}\bar{Q}(\gamma)R^{-1}$$

where we specify $\bar{Q}(\gamma)$ below.

To begin, note that for a fixed vector γ, we can write each row of $g(\gamma, \xi)$ as a quadratic function:

$$g(\gamma, \xi) = \begin{bmatrix} \xi'\widetilde{V}_0\xi & \xi'\widetilde{V}_1\xi & \xi'\widetilde{V}_2\xi & \xi'\widetilde{V}_3\xi & \xi'\widetilde{V}_4\xi & \xi'\widetilde{V}_5\xi \end{bmatrix}'$$

where the $\widetilde{V}_i$ matrices can be written in block form:

$$\widetilde{V}_0 = \begin{bmatrix} 0 & \sigma_1\sigma_2 & 0 \\ 0 & -\sigma_2^2(\gamma_1+\theta) & 0 \\ \sigma_1 h' & \sigma_2(d-h(\gamma_1+\theta))' & h'd \end{bmatrix}$$

$$\widetilde{V}_j = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ \sigma_1 e_j & -\sigma_2(\gamma_1+\theta)e_j & e_j d \end{bmatrix}, \qquad j=1,\dots,5$$

where e_j is a (5×1) vector with a 1 in row j and zeros elsewhere, and

$$d = u[0,0,0,0,1] - \gamma_1 h(\gamma) - \gamma'_{-1}.$$

Note that because of the linearity of our model, we can write the evolution of the vector ξ_n as a (linear) vector autoregression. We then normalize the variable components of ξ to be mean zero, but we retain the constant term. We therefore define the vector $z = \xi - E(\xi) + e_7$. We can then rewrite each row of g as $g^i = z'V_i z + k_i$ for normalized matrices V_i and mean vectors k_i for $i = 0,\dots,5$. Furthermore, we can write the evolution of the vector z as

$$z_{n+1} = Az_n + BW_n$$

for some matrices A and B.

In general, the $\bar{Q}$ matrix depends on the solution of a Poisson equation which can be difficult to solve. However Williams (2001) shows that when the W_i shocks are normally distributed, ξ_n follows a linear VAR, and the g function is quadratic, the (i,j) element of $\bar{Q}$ is given by:

$$\bar{Q}_{i,j} = E\left[(z'L_i z)(z'L_j z) - (z'A'L_iAz + \text{tr}\,(L_iBB'))(z'AL_jAz + \text{tr}\,(L_jBB'))\right], \tag{C.59}$$

where the L_i matrices solve the matrix Lyapunov equations:

$$V_i = L_i - A'L_iA.$$

APPENDIX D. PROOF OF LEMMA 7.2 AND PROPOSITION 7.3

Kushner (1984) pointed out that the large deviation properties of a discrete time recursive stochastic system may not converge to the continuous time approximation of the same discrete time recursive process, unless the associated H-functional of the discrete time process converges to that of the continuous time process. In this application, one can show that the H-functional induced by a multinomial distribution of the perturbation converges as the perturbation converges in distribution to the Gaussian perturbations. Thus, one can approximate the large deviation parameters associated with this recursive system with Gaussian perturbations by another process with a multinomial distribution.

The general characterization of the large deviation parameters presented in the text assumes that W_n has the Gaussian distribution. In order to show how the calculation of the most likely escape point can be derived from the orthogonality condition of the optimal forecasting problem, we shall prove Lemma 7.2 and Proposition 7.3 stated in terms of the multinomial distribution.

To simplify notation, we define a new random variable $\tilde{W}_{in}$ satisfying $\sigma_i W_{in} = \tilde{W}_{in}$: there exist

$$0 < \omega_{i1} < \cdots < \omega_{il} \qquad i = 1,2$$

such that

$$\Pr\left(\tilde{W}_{in} = \omega_{ik}\right) = \Pr\left(\tilde{W}_{in} = -\omega_{ik}\right) = p_{ik} \qquad k = 1,\dots,l \tag{D.60}$$

and

$$\textstyle\sum_{k=1}^{l} p_{ik} = \frac{1}{2} \qquad \forall i \in \{1,2\}.$$

Let

$$\omega_i \in \{-\omega_{il},\dots,-\omega_{i1},\omega_{i1},\dots,\omega_{il}\}$$

be a generic element in the support of $\tilde{W}_{in}$. Let $\omega = (\omega_1, \omega_2)$ be a realization of $\tilde{W}_n$. Since $\tilde{W}_{in}$ $(i = 1, 2)$ takes one of $2l$ different values, $\tilde{W}_n = (\tilde{W}_{1n}, \tilde{W}_{2n})$ can take one of $4l^2$ different values. We use $\tilde{W}_{in}$ in place of $\sigma_i W_{in}$ throughout this proof.

Given $b^1, \dots, b^L \in \Re^2$, let

$$\mathbf{C}(\{b^1, \dots, b^L\})$$

be the cone spanned by $b^1, \dots, b^L$. Given $b \in \Re^2$, let

$$\mathbf{H}(b) = \{x : b \cdot x = 0\}$$

be the hyperplane with directional vector b. Let $\mathbf{H}^+(b)$ be the open half space above $\mathbf{H}(b)$.

Let $\lceil \tau \rceil$ be the integer part of $\tau \in \Re$. Let

$$f^{\varepsilon}_{s,\tau}(\omega) = \frac{\#\{n : \lceil \tau/\varepsilon \rceil \leqq n \leqq \lceil (\tau + s)/\varepsilon \rceil, \tilde{W}_n = \omega\}}{\lceil s/\varepsilon \rceil} \tag{D.61}$$

be the empirical frequency of ω in time interval $[\tau, \tau + s)$. Define

$$f_{s,\tau}(\omega) = \lim_{\varepsilon \to 0} f^{\varepsilon}_{s,\tau}(\omega) \tag{D.62}$$

and

$$f_{\tau}(\omega) = \lim_{s \to 0} f_{s,\tau}(\omega). \tag{D.63}$$

Given that $\tilde{W}_n = (\tilde{W}_{1n}, \tilde{W}_{2n})$ has the multinomial distribution, let $\mathbf{V}$ be the set of all realized values of $\tilde{W}_n$. For each τ, define a subset of $\mathbf{V}$ as

$$\mathbf{V}(\mathbf{H}^+(b)) = \{\omega \in \mathbf{V} : \mathbf{C}(\{R^{-1}\tilde{g}(\gamma(\tau), \omega) : f_{s,\tau}(\omega) > 0\}) \subset \mathbf{H}^+(b)\} \tag{D.64}$$

as the set of perturbations that induce $\{R^{-1}\tilde{g}(\gamma(\tau), \omega)\}$ contained in an open half space.

We state the same half space condition for the multinomial distribution.

Definition D.1. *If $\omega, \omega' \in \mathbf{V}(\mathbf{H}^+(b))$, then ω and ω' satisfy the same half space condition at τ. We say that γ satisfies the same half space condition, if there exists $b \in \Re^2$ such that $\gamma(\tau)$ is generated by perturbations in $\mathbf{V}(\mathbf{H}^+(b))$ for all τ.*

Note that $\mathbf{V}$ contains exactly $4l^2$ elements and each hyperplane divides $\mathbf{V}$ into 2 subsets, each of which has $2l^2$ elements. For $\forall \tau > 0$, we can choose $b^1, \dots, b^{2l^2}$ such that

$$\mathbf{H}_k = \begin{cases} \mathbf{H}(b^k) & \text{if } k \leqq 2l^2 \\ \mathbf{H}(-b^{k-2l^2}) & \text{if } k > 2l^2 \end{cases}$$

and

$$\bigcup_{b \in \Re^2} \mathbf{V}(\mathbf{H}(b)) = \bigcup_{k=1}^{4l^2} \mathbf{V}(\mathbf{H}_k)$$

and each $\mathbf{V}(\mathbf{H}_k)$ contains $2l^2$ elements.

Lemma D.2. *If γ^e is a dominant escape path with exit time $\tau^e < \infty$, then the same half-space condition must hold for all $\tau \in (0, \tau^e)$ along the trajectory of γ^e. Moreover, there exists $k \in \{1, \dots, 4l^2\}$ such that*

$$\dot{\gamma}^e(\tau) = \sum_{\omega} R^{-1}\tilde{g}(\gamma^e(\tau), \omega) \Pr(\omega : \mathbf{V}(\mathbf{H}_k)) \qquad \forall \tau \in (0, \tau^e). \tag{D.65}$$

If γ^{e,k_1} is the trajectory induced by (D.65) *that has the shortest escape time among $4l^2$ escape paths induced by* (D.65), *then γ^{e,k_1} is a dominant escape path.*

Proof. Fix $\mathbf{V}(\mathbf{H}_k)$ and calculate the "conditional" ODE $\gamma^{e,k}$ on $\mathbf{V}(\mathbf{H}_k)$:

$$\gamma^{e,k}(\tau) = \gamma^{e,k}(0) + \lim_{\varepsilon \to 0} \varepsilon \sum_{t=1}^{\lceil \tau/\varepsilon \rceil} R^{-1}\tilde{g}(\gamma_n, \omega_n)$$

where $\omega_n \in \mathbf{V}(\mathbf{H}_k)$ for $\forall n = 1, \dots, \lceil \tau/a \rceil$ and the probability distribution over $\mathbf{V}(\mathbf{H}_k)$ is the conditional probability distribution of $\mathbf{V}$ on $\mathbf{V}(\mathbf{H}_k)$. By the construction of $G(\tau)$,

$$\gamma^{e,k}(\tau) = x^{e,k} \in \partial G(\tau).$$

Fix $\gamma^{e,k}$, $\rho > 0$, and τ^e. Define

$$N_\rho(\gamma^{e,k}) = \{\gamma \in \Re^2 : \exists \tau < \tau^e, |\gamma - \gamma^{e,k}(\tau)| < \rho\}$$

as the ρ neighbourhood of the trajectory of $\gamma^{e,k}$ between 0 and τ^e. Define

$$A_\varepsilon(\mathbf{H}_k) = \{\gamma^\varepsilon : \gamma^\varepsilon(0) = \gamma^{e,k}, \exists \tau < \tau^e, \gamma^\varepsilon(\tau) \notin N_\rho(\gamma^{e,k})\}$$

be the set of sample paths generated by perturbations in $\mathbf{V}(\mathbf{H}_k, \tau)$, which move away from the neighbourhood of the trajectory sometime between 0 and τ^e.

The next result is an extension of Cramér's theorem to the probability distributions over function spaces, whose proof can be found in Freidlin (1978) and Dupuis and Kushner (1985) under general conditions.

Proposition D.3. *For* $\forall \rho > 0$, *there exists* $s^* > 0$ *such that*

$$\lim_{\varepsilon \to 0} \varepsilon \log \Pr(A_\varepsilon(\mathbf{H}_k) : \mathbf{V}(\mathbf{H}_k)) \leqq -s^*.$$

Fix τ^e and define

$$x^{e,k} = \gamma^{e,k}(\tau^e)$$

for each $k = 1, \ldots, 4l^2$. Define $G(\tau^e)$ as the convex hull of

$$X^e = \{x^{e,1}, \ldots, x^{e,4l^2}\}.$$

By the construction, $G(\tau^e)$ is a convex polyhedra with extreme points selected from X^e. Clearly, if $\tau > \tau^e$, then

$$G(\tau^e) \subset G(\tau).$$

Since $G(\tau^e)$ expands as τ^e increases, we can find the first time when some element of X^e crosses over ∂G:

$$\tau^{e,1} = \inf \{\tau : X^e \setminus G \neq \emptyset\}.$$

If $X^e \cap \partial G$ contains more than a single element, we can enumerate them in an arbitrary manner. Let x^{e,k_1} be the element in $X^e \cap \partial G$ when $\tau^e = \tau^{e,1}$.

Choose $\rho > 0$ sufficiently small such that

$$N_\rho(x^{e,k}) \cap N_\rho(x^{e,k'}) = \emptyset \qquad \forall k \neq k'.$$

Since $\partial G(\tau^e) \setminus N_\rho(X^e)$ is compact, we can find $y_1, \ldots, y_L \in \partial G(\tau^e) \setminus N_\rho(X^e)$ and $\rho_1, \ldots, \rho_L > 0$ such that

$$\partial G(\tau^e) \setminus N_\rho(X^e) \subset \bigcup_{k=1}^{L} N_{\rho_k}(y_k).$$

It remains to prove that x^{e,k_1} is the most likely exit point through G. To this end, we choose G a convex compact neighbourhood of the stable point $\bar{\gamma}$, and choose $\tau^e = \tau^{e,1}$ as defined above. Define $G(\tau^e)$ as the convex hull of X^e. By the construction,

$$G(\tau^e) \cap \partial G = \{x^{e,k_1}\}.$$

In order to analyse the probability distribution over G induced by the escape paths, it is necessary to "count" the escape path through $N_\rho(y) \cap \partial G$ for each $y \in \partial G$ and $\rho > 0$.

Divide the "arc" $N_\rho(y) \cap G$ into smaller pieces, each of which is as long as ε and centred around y_k. There are approximately as many as $\lceil \rho/\varepsilon \rceil$ such small grids contained in $N_\rho(y) \cap \partial G$. For convenience, we abuse the notation by writing the centres of the small grids as $y_1, \ldots, y_{\lceil \rho/\varepsilon \rceil}$, and $N_\varepsilon(y_k)$ as the small grid centred around y_k ($k = 1, \ldots, \lceil \rho/\varepsilon \rceil$). For each grid $N_\varepsilon(y_k)$, we can find the shortest escape path through $N_\varepsilon(y_k)$. Let $A_0^k(y, \rho, \varepsilon)$ be the set of all shortest paths through $N_\varepsilon(y_k)$. Define

$$A_0(y, \rho, \varepsilon) = \bigcup_{k=1}^{\lceil \rho/\varepsilon \rceil} A_0^k(y, \rho, \varepsilon).$$

For $j \geqq 1$, let $A_j^k(y, \rho, \varepsilon)$ be the set of escape paths through $N_\varepsilon(y_k)$ that takes j periods more than the shortest paths through $N_\varepsilon(y_k)$. Define

$$A_j(y, \rho, \varepsilon) = \bigcup_{k=1}^{\lceil \rho/\varepsilon \rceil} A_j^k(y, \rho, \varepsilon).$$

Note that $A_j^k(y, \rho, \varepsilon)$ increases as $\varepsilon \to 0$ for a fixed G. But, the rate of increase is asymptotically bounded by ε^{-j}. Define

$$\alpha_j(y, \rho, \varepsilon) = \frac{\#A_j(y, \rho, \varepsilon)}{\#A_0(y, \rho, \varepsilon)}$$

where $\#A$ is the number of elements in A. Again, $\alpha_j(y, \rho, \varepsilon)$ increases as $\varepsilon \to 0$ asymptotically at the rate of ε^{-j}. Let $\mathbf{P}(y, \rho, \varepsilon)$ be the probability of escape through $N_\rho(y) \cap \partial G$.

As $\varepsilon \to 0$, each element in $A_0(y, \rho, \varepsilon)$ is the shortest path around a very small neighbourhood of y_k $(k = 1, \ldots, \lceil \rho/\varepsilon \rceil)$. For this reason, we call each element in $A_0(y, \rho, \varepsilon)$ the *pointwise dominant* path. Define

$$p(y, \rho, \varepsilon) = \Pr(A_0(y, \rho, \varepsilon))$$

as the probability of escaping through the pointwise dominant paths.

The next proposition shows that any pointwise dominant path must be generated by perturbations that satisfy the same half-space condition along its trajectory.

Proposition D.4. *As $\varepsilon \to 0$, almost all escape paths in $A_0(y, \rho, \varepsilon)$ must satisfy the same half-space condition.*

Proof of Proposition D.4. If the conclusion of Proposition D.4 is false, then with a positive probability, we can find an escape path γ with

$$\gamma(\tau^e) \in \partial G(\tau^e),$$

$\tau \in (0, \tau^e)$ and $\{s_k\}_{k=1}^{\infty}$ such that

$$s_k \downarrow 0$$

and $\forall s_k > 0$,

$$\mathbf{C}(\{R^{-1}\tilde{g}(\gamma(\tau), \omega) : f_{s_k,\tau}(\omega) > 0\}) = \Re^2. \quad || \tag{D.66}$$

Because the next lemma is a straightforward application of linear algebra, we state the result without the proof.

Lemma D.5. (D.66) *holds, if and only if there exist ω^1, ω^2 and ω^- in the support of $f_{s_k,\tau}$ such that*

$$R^{-1}\tilde{g}(\gamma(\tau), \omega^-) \in \mathbf{C}\big(\{-R^{-1}\tilde{g}(\gamma(\tau), \omega^1), -R^{-1}\tilde{g}(\gamma(\tau), \omega^2)\}\big). \tag{D.67}$$

If ω^1, ω^2 and ω^- satisfy (D.67), *then any permutation of the three vectors satisfies* (D.67).

Proof of Lemma D.5. See Cho and Sargent (1999) ||

Fix $s = s_k > 0$. Let

$$\textstyle\sum_\omega R^{-1}\tilde{g}(\gamma(\tau), \omega) f_{s,\tau}(\omega)$$

be the mean directional vector at τ. By Lemma D.5, we can choose ω^1, ω^2 and ω^- such that

$$\mathbf{C}\big(\{R^{-1}\tilde{g}(\gamma(\tau), \omega^1), R^{-1}\tilde{g}(\gamma(\tau), \omega^2), R^{-1}\tilde{g}(\gamma(\tau), \omega^3)\}\big) = \Re^2.$$

Rename each vector so that

$$\textstyle\sum_\omega R^{-1}\tilde{g}(\gamma(\tau), \omega) f_{s,\tau}(\omega) \in \mathbf{C}\big(\{R^{-1}\tilde{g}(\gamma(\tau), \omega^1), R^{-1}\tilde{g}(\gamma(\tau), \omega^2)\}\big) \tag{D.68}$$

and let $\omega^3 = \omega^-$ so that there exist $\rho^1(\omega^-) > 0$ and $\rho^2(\omega^-) > 0$ such that

$$R^{-1}\tilde{g}(\gamma(\tau), \omega^-) = -\rho^1(\omega^-)R^{-1}\tilde{g}(\gamma(\tau), \omega^1) - \rho^2(\omega^-)R^{-1}\tilde{g}(\gamma(\tau), \omega^2).$$

To simplify notation, let us assume that there is only one such ω^- in the support of $f_{s,\tau}$. The general case follows from exactly the same logic, while the notation becomes significantly more complicated.

Note that for $\omega \in [0, 1]$,

$$\begin{aligned}\textstyle\sum_\omega R^{-1}\tilde{g}(\gamma(\tau), \omega) f_{s,\tau}(\omega) &= \textstyle\sum_\omega R^{-1}\tilde{g}(\gamma(\tau), \omega) f_{s,\tau}(\omega) + O(s)\\ &= \textstyle\sum_{j=1}^{2}\big[\sum_{\omega \neq \omega^-} \rho^j(\omega) f_{s,\tau}(\omega) + \omega\rho^j(\omega^-) f_{s,\tau}(\omega^-)\big] R^{-1}\tilde{g}(\gamma(\tau), \omega^j)\\ &\quad + (1 - \omega) f_{s,\tau}(\omega^-) R^{-1}\tilde{g}(\gamma(\tau), \omega^-) + O(s)\end{aligned}$$

Since

$$\left[\sum_{\omega\neq\omega^-}\rho^j(\omega)f_{s,\tau}(\omega)\right]+\rho^j(\omega^-)f_{s,\tau}(\omega^-)>0 \quad \forall j=1,2,$$

there exists $h^j\in(0,1)$ such that

$$h^j\left[\sum_{\omega\neq\omega^-}\rho^j(\omega)f_{s,\tau}(\omega)\right]+\rho^j(\omega^-)f_{s,\tau}(\omega^-)=0$$

which implies that

$$f_{s,\tau}(\omega^-)=-\frac{h^j}{\rho^j(\omega^-)}\sum_{\omega\neq\omega^-}\rho^j(\omega)f_{s,\tau}(\omega).$$

Without loss of generality, assume that

$$-\frac{1}{\rho^1(\omega^-)}\sum_{\omega\neq\omega^-}\rho^1(\omega)f_{s,\tau}(\omega)<-\frac{1}{\rho^2(\omega^-)}\sum_{\omega\neq\omega^-}\rho^2(\omega)f_{s,\tau}(\omega),$$

which implies that

$$h^1>h^2.$$

Hence,

$$h^1\left[\sum_{\omega\neq\omega^-}\rho^2(\omega)f_{s,\tau}(\omega)\right]+\rho^2(\omega^-)f_{s,\tau}(\omega^-)>0.$$

We can write the mean dynamics as

$$\begin{aligned}\sum_\omega R^{-1}\tilde{g}(\gamma(\tau),\omega)f_{s,\tau}(\omega)&=\left[(1-\omega h^1)\sum_{\omega\neq\omega^-,\omega^2}R^{-1}\tilde{g}(\gamma(\tau),\omega)f_{s,\tau}(\omega)\right]\\&\quad+\left[(1-\omega h^1)f_{s,\tau}(\omega^2)+\rho^2(\omega^-)\omega(h^1-h^2)f_{s,\tau}(\omega^-)\right]R^{-1}\tilde{g}(\gamma(\tau),\omega^2)\\&\quad+(1-\omega)f_{s,\tau}(\omega^-)R^{-1}\tilde{g}(\gamma(\tau),\omega^-)+O(s).\end{aligned}$$

Note that

$$\begin{pmatrix}\left[(1-\omega h^1)\Sigma_{\omega\neq\omega^-,\omega^2}f_{s,\tau}(\omega)\right]+\left[(1-\omega)f_{s,\tau}(\omega^-)\right]\\ \qquad+\left[(1-\omega h^1)f_{s,\tau}(\omega^2)+\rho^2(\omega^-)\omega(h^1-h^2)f_{s,\tau}(\omega^-)\right]\end{pmatrix}<1$$

and that each term inside the bracket is strictly positive for a sufficiently small $\omega\in(0,1)$. We can generate the same mean dynamics through

$$\tilde{f}(\omega)=\begin{cases}(1-h^1)f_{s,\tau}(\omega) & \text{if } \omega\neq\omega^-,\omega^2\\ (1-\omega h^1)f_{s,\tau}(\omega^2)+\rho^2(\omega^-)\omega(h^1-h^2)f_{s,\tau}(\omega^-) & \text{if } \omega=\omega^2\\ (1-\omega)f_{s,\tau}(\omega^-) & \text{if } \omega=\omega^-\end{cases} \tag{D.69}$$

instead of $f_{s,\tau}$. Since $\Sigma_\omega\tilde{f}(\omega)<1$, $\tilde{f}$ is not an empirical frequency. After normalization, we have

$$\gamma(\tau+s)-\gamma(\tau)=s\left[\sum_\omega\tilde{f}(\omega)+O(s)\right]\left[\sum_\omega\frac{\tilde{f}(\omega)}{\Sigma_\omega\tilde{f}(\omega)}R^{-1}\tilde{g}(\gamma(\tau),\omega)\right]+sO(s). \tag{D.70}$$

For a sufficiently small $s=s_k>0$,

$$\sum_\omega\tilde{f}(\omega)+O(s)<1$$

which implies that we can construct a path from $\gamma(\tau)$ to $\gamma(\tau+s)$ which takes $s[\Sigma_\omega\tilde{f}(\omega)+O(s)]$ instead of s. This contracts the hypothesis that γ is the pointwise dominant path to $x^e\in\partial G$. ||

For each $\omega^k\in\mathbf{V}$, let

$$p^k=\Pr\left(\omega^k\right)>0$$

be the probability of the k-th perturbation. Define

$$\bar{p}=\max\left(p^1,\ldots,p^{4l^2}\right)$$

and

$$\underline{p} = \min\left(p^1, \ldots, p^{4l^2}\right).$$

Then,

$$\sum_{j=0}^{\infty} p(y, \rho, \varepsilon)\alpha_j(y, \rho, a)\underline{p}^j \leqq \mathbf{P}(y, \rho, \varepsilon) \leqq \sum_{j=0}^{\infty} p(y, \rho, \varepsilon)\alpha_j(y, \rho, \varepsilon)\bar{p}^j.$$

Note that

$$\frac{\sum_{k=1}^{L} p(y_k, \rho_k, \varepsilon)\sum_{j=0}^{\infty}\alpha_j(y_k, \rho_k, \varepsilon)\bar{p}^j}{\sum_{k=1}^{4l^2} p(x_k, \rho, \varepsilon)\sum_{j=0}^{\infty}\alpha_j(x^{e,k}, \rho, \varepsilon)\underline{p}^j} \leqq \sum_{k=1}^{L}\sum_{j=0}^{\infty}\frac{p(y_k, \rho_k, \varepsilon)\sum_{j=0}^{\infty}\alpha_j(y_k, \rho_k, \varepsilon)\bar{p}^j}{\sum_{k=1}^{4l^2} p(x^{e,k}, \rho, \varepsilon)}. \tag{D.71}$$

Recall that $G(\tau^e)(\subset G)$ is the convex hull of the points that can be reached in τ^e by $4l^2$ escape paths. Thus, it takes an equal time to reach the boundary of $G(\tau^e)$ as long as the escape path satisfies the same half space condition. Since $y_k \notin \bigcup_{k=1}^{4l^2} N_\rho(x^{e,k})$, any escape path γ through $N_{\rho_k}(y_k)$ must be away from $\gamma^{e,k}$ by at least as much as ρ: $\exists\tau' \leqq \tau^e$,

$$\gamma(\tau') \notin N_\rho(\gamma^{e,k}) \quad \forall k.$$

Thus,

$$p(y_k, \rho_k, \varepsilon) \leqq \Pr\left(\bigcap_k \{\gamma : \exists\tau' \leqq \tau^e, \gamma(\tau') \notin N_\rho(\gamma^{e,k})\}\right)$$
$$\leqq \Pr\left(\{\gamma : \exists\tau' \leqq \tau^e, \gamma(\tau') \notin N_\rho(\gamma^{e,1})\}\right).$$

Thus, Proposition D.3 implies that there exists $s^* > 0$ such that

$$\frac{p(y_k, \rho_k, \varepsilon)}{\sum_{k=1}^{4l^2} p(x^{e,k}, \rho, \varepsilon)} \leqq \frac{p(y_k, \rho_k, \varepsilon)}{p(x^{e,1}, \rho, \varepsilon)} \leqq e^{-s^*/\varepsilon}$$

for any small $a > 0$. Thus, (D.71) can be bounded by

$$\sum_{k=1}^{L}\sum_{j=0}^{\infty} e^{-s^*/\varepsilon}\alpha_j(y_k, \rho_k, \varepsilon)\bar{p}^j.$$

Let

$$f_a(j : k) = e^{-s^*/\varepsilon}\alpha_j(y_k, \rho_k, \varepsilon)$$

which converges to 0 pointwise as $\varepsilon \to 0$. Thus,

$$\limsup_{\varepsilon\to 0} f_\varepsilon(j : k) = 0 \quad \forall j, k.$$

Since

$$\sum_{j=0}^{\infty} \bar{p}^j < \infty,$$

$\{f_\varepsilon(j : k)\}_a$ is a sequence of integrable functions with respect to the finite measure μ where

$$\mu(j) = \bar{p}^j \quad \forall j.$$

Hence,

$$0 \leqq \limsup_{\varepsilon\to 0} \sum_{k=1}^{L}\sum_{j=0}^{\infty} e^{-s^*/\varepsilon}\alpha_j(y_k, \rho_k, \varepsilon)\bar{p}^j \leqq \sum_{k=1}^{L}\sum_{j=0}^{\infty}\limsup_{\varepsilon\to 0}\left[e^{-s^*/\varepsilon}\alpha_j(y_k, \rho_k, \varepsilon)\right]\bar{p}^j = 0$$

which implies that the probability distribution over $G(\tau^e)$ induced by the escape points is concentrated at $N_\rho(X^e)$. Since the only intersection between X^e and ∂G is x^{e,k_1}, this result proves that x^{e,k_1} is the most likely escape point within $\tau^{e,1}$ real time.

For $\tau^e > \tau^{e,1}$, the same logic applies. Proposition D.3 implies that the probability distribution induced by the exit points is concentrated at the intersection of the trajectory of $\gamma^{e,k}$ and ∂G. Note that there are as many as $4l^2$ such trajectories. Among $4l^2$ exit points of the trajectories induced by the "conditional" ODE, it takes the shortest time to escape through x^{e,k_1} along the trajectory γ^{e,k_1} induced by the conditional ODE. Since the longer escape paths requires exponentially more "unusual" perturbations, the probability of escaping through the shorter path is exponentially higher than that through the longer path. Thus, the exit probability is concentrated at $x^{e,1}$ which is the shortest escape paths among the escape paths induced by the conditional ODEs. ||

APPENDIX E. NUMERICAL SOLUTION

As mentioned in Section 6, in order to calculate the dominant escape path numerically, it helps to recast the problem as an initial value problem rather than a two-point boundary value problem. For any arbitrary $(t, \hat{\gamma})$ there may not be a solution to (5.43) such that $\gamma(t) = \hat{\gamma}$. However, given initial conditions $(a(0), \lambda(0)) = \alpha_0$, we can determine candidate values of t and $\hat{\gamma}$ endogenously by integrating (5.43) until the system leaves G, which gives $t = \inf\{s : \gamma(s) \notin G\}$, and $\hat{\gamma} = \gamma(t)$. So rather than attempting to directly calculate $\bar{S}$ by minimizing over $(t, \hat{\gamma})$, which may not be solvable numerically, we can instead define:

$$S(\alpha_0) = \int_0^t L(\gamma(s), \dot{\gamma}(s); R(s))ds, \quad \text{s.t. } (5.43), \gamma(0) = \bar{\gamma}, R(0) = \bar{R}, (a(0), \lambda(0)) = \alpha_0 \tag{E.72}$$

with t defined as in the previous sentence. Minimizing this quantity over the vector α_0 will then give us $\bar{S}$ and an endogenously determined escape point $\hat{\gamma}$ and the minimizing path is the dominant escape path.

In order to compute the dominant escape path, we solve the ODEs using Runge–Kutta methods, and minimize over α_0 numerically, where the integrals in S are also calculated numerically with a trapezoidal rule. We also compute the derivatives in (5.43) using the central finite differences:

$$\frac{\partial \bar{g}}{\partial \gamma_i} = \frac{\bar{g}(\gamma + e_i\Delta) - \bar{g}(\gamma - e_i\Delta)}{2\Delta}, \quad \frac{\partial Q}{\partial \gamma_i} = \frac{Q(\gamma + e_i\Delta, R) - Q(\gamma - e_i\Delta, R)}{2\Delta}, \tag{E.73}$$

for some small $\Delta > 0$. The dependence of Q and $\bar{g}$ on γ is complex, as both the conditional and invariant distributions change with γ, which is why we compute the derivatives numerically. It is straightforward but tedious to derive the exact expressions for the derivatives, which can be calculated explicitly up to the solution of matrix Lyapunov equations. In practice, the finite difference method proved slightly faster and only very slightly less accurate than the explicit calculations.

Acknowledgements. We thank Xiahong Chen, Amir Dembo, George Evans, Lars Peter Hansen, Michael Harrison, Seppo Honkapohja, Peter Howitt, Kenneth Judd, Robert King, David M. Kreps, and Eric Swanson for helpful discussions, and three anonymous referees for insightful criticisms that helped us to improve the substance and exposition of the paper. Chao Wei helped with the simulations. Cho and Sargent gratefully thank the National Science Foundation for Research support. Williams's research was supported by a fellowship in Applied Economics from the Social Science Research Council, with funds provided by the John D. and Catherine T. MacArthur Foundation. Sargent chose the order of authors' names to acknowledge Williams's special contributions in cracking difficult technical aspects of the problem.

REFERENCES

BARRO, R. J. and GORDON, D. B. (1983), "Rules, Discretion, and Reputation in a Model of Monetary Policy", *Journal of Monetary Economics*, **12**, 101–121.

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

BULLARD, J. and CHO, I.-K. (2001), "Escapist Policy Rules" (mimeo).

CHO, I.-K. and MATSUI, A. (1995), "Induction and the Ramsey Policy", *Journal of Economic Dynamics and Control*, **19**, 1113–1140.

CHO, I.-K. and SARGENT, T. J. (1999), "Escaping Nash Inflation" (mimeo).

CHUNG, H. (1990) *Did Policy Makers Really Believe in the Phillips Curve?: An Econometric Test* (Ph.D. thesis, University of Minnesota).

DEMBO, A. and ZEITOUNI, O. (1990) *Large Deviations Techniques and Applications*, 2nd ed. (New York: Springer-Verlag).

DUPUIS, P. and KUSHNER, H. J. (1985), "Stochastic Approximation via Large Deviations: Asymptotic Properties", *SIAM Journal of Control and Optimization*, **23**, 675–696.

DUPUIS, P. and KUSHNER, H. J. (1989), "Stochastic Approximation and Large Deviations: Upper Bounds and w.p.1 Convergence", *SIAM Journal of Control and Optimization*, **27**, 1108–1135.

EVANS, G. W. and HONKAPOHJA, S. (1993), "Adaptive Forecasts, Hysteresis and Endogenous Fluctuations", *Economic Review of Federal Reserve Bank of San Francisco*, **1**, 3–13.

EVANS, G. W. and HONKAPOHJA, S. (1998), "Economic Dynamics with Learning: New Stability Results", *Review of Economic Studies*, **65**, 23–44.
EVANS, G. W. and HONKAPOHJA, S. (2001) *Learning and Expectations in Macroeconomics* (Princeton University Press).
FREIDLIN, M. I. (1978), "The Average Principle and Theorems on Large Deviations", *Russian Mathematical Surveys*, **33**, 117–176.
FREIDLIN, M. I. and WENTZELL, A. D. (1984) *Random Perturbations of Dynamical Systems* (Springer-Verlag).
FUDENBERG, D. and KREPS, D. M. (1995), "Learning in Extensive Games, I: Self-Confirming and Nash Equilibrium", *Games and Economic Behavior*, **8**, 20–55.
FUDENBERG, D. and LEVINE, D. (1998) *Learning in Games* (Cambridge: MIT Press).
FUDENBERG, D. and LEVINE, D. K. (1993), "Self-Confirming Equilibrium", *Econometrica*, **61**, 523–545.
KANDORI, M., MAILATH, G. and ROB, R. (1993), "Learning, Mutation and Long Run Equilibria in Games", *Econometrica*, **61**, 27–56.
KASA, K. (1999), "Learning, Large Deviations and Recurrent Currency Crises" (mimeo, Federal Reserve Bank of San Francisco).
KREPS, D. M. (1998), "Anticipated Utility and Dynamic Choice", in *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983–1997* (Cambridge: Cambridge University Press).
KUSHNER, H. J. (1984), "Robustness and Approximation of Escape Time and Large Deviation Estimates for Systems with Small Noise Effects", *SIAM Journal of Applied Mathematics*, **44**, 160–182.
KUSHNER, H. J. and YIN, G. G. (1997) *Stochastic Approximation Algorithms and Applications* (Berlin: Springer-Verlag).
KYDLAND, F. and PRESCOTT, E. C. (1977), "Rules Rather than Discretion: the Inconsistency of Optimal Plans", *Journal of Political Economy*, **85**, 473–491.
MARCET, A. and NICOLINI, J. P. (1997), "Recurrent Hyperinflations and Learning" (mimeo, Universitat Pompeu Fabra).
MARCET, A. and SARGENT, T. J. (1989*b*), "Convergence of Least Squares Learning in Environments with Hidden State Variables and Private Information", *Journal of Political Economy*, **97**, 1306–1322.
MARCET, A. and SARGENT, T. J. (1989*a*), "Convergence of Least Squares Learning Mechanisms in Self Referential Linear Stochastic Models", *Journal of Economic Theory*, **48**, 337–368.
PHELPS, E. S. (1967), "Phillips Curves, Expectations of Inflation and Optimal Unemployment Over Time", *Economica*, **2**, 22–44.
SARGENT, T. J. (1999) *The Conquest of American Inflation* (Princeton: Princeton University Press).
SIMS, C. A. (1972), "The Role of Approximate Prior Restrictions on Distributed Lag Estimation", *Journal of the American Statistican Association*, **67**, 169–175.
SIMS, C. A. (1988), "Projecting Policy Effects with Statistical Models", *Revista de Analisis Economico*, **3**, 3–20.
STOKEY, N. L. (1989), "Reputation and Time Consistency", *American Economic Review, Papers and Proceedings*, **79**, 134–139.
WHITE, H. (1982), "Maximum Likelihood Estimation of Misspecified Models", *Econometrica*, **50**, 1–25.
WHITTLE, P. (1996) *Optimal Control: Basics and Beyond* (New York: Wiley).
WILLIAMS, N. (2001) *Escape Dynamics in Learning Models* (Ph.D. thesis, University of Chicago).
WOODFORD, M. D. (1990), "Learning to Believe in Sunspots". *Econometrica*, **58**, 277–307.
WORMS, J. (1999), "Moderate Deviations for Stable Markov Chains and Regression Models", *Electronic Journal of Probability*, **4** 1–28.
YOUNG, P. (1993), "The Evolution of Conventions", *Econometrica*, **61**, 57–83.