# Escape Dynamics in Learning Models

Noah Williams*

*Department of Economics, University of Wisconsin - Madison*
E-mail: nmwilliams@wisc.edu

Revised January 28, 2009

Learning models have typically focused on the convergence of beliefs toward an equilibrium. However in stochastic environments, there may be rare but recurrent episodes where shocks cause beliefs to escape from the equilibrium. We characterize these escape dynamics by drawing on the theory of large deviations, developing new results which make this theory directly applicable in a class of linear-quadratic learning models. The likelihood, frequency, and most likely direction of escapes are all characterized by a simple deterministic control problem. We illustrate our results in a simple example, which shows how escapes can arise in a model with a unique equilibrium.

## 1. INTRODUCTION

There is by now a substantial literature on adaptive learning in economics, with important work both in macroeconomics (see Evans and Honkapojha, 2001) and game theory (see Fudenberg and Levine, 1998). The central question in this literature has been whether learning leads to equilibrium behavior. As agents who use simple learning rules observe more and more data, their beliefs may settle down and converge to an equilibrium. Large movements away from this equilibrium then become increasingly unlikely, but due to ongoing stochastic shocks they may occasionally occur. In this paper we develop and apply methods to characterize these rare departures. Following Sargent (1999), we call these large deviations *escape dynamics*.

In particular, we focus on situations in which agents are uncertain of their economic environment. These agents base their actions on subjective models, which we take to be linear regressions that they update as they observe data. Agents allow for structural change in their environment, which leads them to discount past data. We also consider the possibility that agents' subjective models may be misspecified, and thus following Sargent (1999) our equilibrium concept is a self-confirming equilibrium (SCE).[1] In an SCE, agents' beliefs are correct about outcomes that occur with positive probability, but may be incorrect about events which happen with probability zero.

*I thank Andrew Postlewaite (the editor) and two anonymous referees for comments that greatly improved the paper. I also thank Fernando Alvarez, Jim Bullard, Marco Cagetti, Jeff Campbell, Xiaohong Chen, In-Koo Cho, Amir Dembo, John Duffy, Dana Heller, Ken Kasa, Juha Seppälä, Ted Temzelides, Harald Uhlig, and especially Lars Peter Hansen and Thomas Sargent for helpful comments, discussions, and suggestions. This paper is based on my Ph.D. dissertation at the University of Chicago. I am grateful for the support that I received during that time through a fellowship in Applied Economics from the Social Science Research Council, with funds provided by the John D. and Catherine T. MacArthur Foundation. This revised version is based upon work supported by the National Science Foundation under Grant No. 0317848.

[1] See Fudenberg and Levine (1998) for further discussion and background on this equilibrium concept.

Since agents discount past data, their beliefs do not converge as they obtain more observations. Thus we study an alternate limit in which the discount rate on past data, known as the *gain*, gets small. With small gains, agents average more evenly over past data and a law of large numbers applies. We show that agents' beliefs converge to a differential equation, which we call the *mean dynamics*. On average, the mean dynamics pull agents toward a self-confirming equilibrium. This result parallels much of the adaptive learning literature, beginning with Marcet and Sargent (1989). Our specific result applies stochastic approximation theory due to Kushner and Yin (1997), and is analogous to results in Evans and Honkapohja (2001).

Occasionally however, an accumulation of stochastic shocks may induce agents to change their beliefs, and so they may escape the self-confirming equilibrium. In the limit, the probability of escaping from the SCE goes to zero, and thus escapes become increasingly unlikely for smaller gain settings. But for any positive gain, escapes may occasionally recur, and when they do, they have a very particular form. To analyze the escape dynamics, we draw on the theory of large deviations. We show that when agents' beliefs escape, with high probability they closely follow a deterministic path we call the *most probable escape path*. We show how to find this path, and we characterize the likelihood and frequency of escapes.

Our main contribution is the derivation of a simple deterministic control problem whose solution characterizes the escape dynamics in a class of linear-quadratic models. Our results build on the general analysis of Dupuis and Kushner (1989). They develop a theory of large deviations for stochastic approximation models, and characterize escapes via a variational problem.[2] While their results are quite general, they are difficult to apply in practice. We provide new results which simplify the general theory in our particular linear-quadratic setting. Drawing on a result from Worms (1999), we reformulate the general variational problem as a simple cost minimization problem.

In the minimization problem, we clearly separate the two sources of dynamics which govern agents' beliefs. The mean dynamics govern the expected behavior of beliefs and drive the convergence results. Escape dynamics are driven by unlikely shock realizations, and we show that they can be interpreted as a perturbation of the mean dynamics. Our key results derive a cost function which provides a measure of the likelihood of the perturbations. The most probable escape path can be found by choosing a minimum cost sequence of perturbations which push agents' beliefs away from the SCE. We then apply standard control theory methods to characterize the solution of the cost minimization problem.

While the mean dynamics and convergence to an equilibrium have been well-studied in the literature, there has been much less focus on escape dynamics. The insight that stochastic shocks may push agents away from an equilibrium has been most extensively analyzed in evolutionary game theory.[3] This literature has focused on games with multiple equilibria, and used large deviation methods to determine the stochas-

[2]Dupuis and Kushner (1989) in turn build on Freidlin and Wentzell (1998), who developed large deviations for continuous time diffusion processes.

[3]Important papers in this literature include Foster and Young (1990), Kandori, Mailath, and Rob (1993) and Young (1993). See Section 4.4 below for more discussion. There are also some technical differences, as

tic transition rates between equilibria. Although our results can be used to analyze multiplicities as well, in this paper we focus on models with a unique equilibrium.

While our methods can be applied to a variety of models, not all of them will have prominent escape dynamics. In many models large deviations from an equilibrium would be quite infrequent, and our results would characterize rare tail events. However, Sargent (1999) demonstrated that escape dynamics may be an important force in some models with a unique equilibrium. He also introduced large deviation theory for settings like ours, and provided much of the motivation for this paper. Since the early drafts of this paper have circulated, the results developed here have been applied in a variety of settings. Our results were first applied by Cho, Williams, and Sargent (2002) to analyze Sargent's (1999) model. Further related papers which build on or apply our results include Sargent and Williams (2005), Bullard and Cho (2005), McGough (2006), Ellison and Yates (2007), Cho and Kasa (2008), and Williams (2009). See Section 4.4 for more discussion.

To illustrate our methods, we develop a simple example which shares many features of the more elaborate economic models considered in the literature. We study a model of a monopolist learning a demand curve, subject to both cost and demand shocks. We show that the model has a unique self-confirming equilibrium, and the firm's beliefs converge to it in the small gain limit. However when the cost and demand shocks are correlated, there are recurrent episodes in which the firm's beliefs escape from the SCE, and the firm rapidly raises its price. After such an escape, the beliefs are gradually drawn back to the SCE and the price falls back to the SCE level. These escapes lead to recurrent large price fluctuations, which we do not observe when the cost and demand shocks are independent. We apply our results to explain this difference, and to characterize both the frequency of escapes and the behavior of the firm's beliefs during an escape. We then describe how the escapes are due to locally self-reinforcing dynamics which kick in once stochastic shocks push beliefs away from the SCE.

The rest of the paper is organized as follows. In Section 2 we introduce our baseline model, a single-agent linear quadratic model, and discuss our equilibrium concept and learning formulation. In Section 3 we provide an overview of our results. We first briefly describe the convergence of beliefs, and then turn to the analysis of escape dynamics, concluding with our main result which characterizes them. Section 4 then describes and analyzes an example. The final sections of the paper provide more formal detail for the results in Section 3. In Section 5, we formally establish the convergence of beliefs, and Section 6 provides the large deviation results which characterize the escape dynamics. The appendix collects proofs and statements of technical results.

## 2. THE MODEL

In this section we describe the class of models we study in the paper. We focus on linear models in which agents form decisions based on estimated models which they update over time. For simplicity, we focus on a single agent setting. Our results can

the game theory literature has generally focussed on models with finite state spaces, unlike the continuous state space we study here.

also be applied to dynamic games, as in Williams (2009), but strategic interaction raises some additional issues which are not essential here.

### 2.1. The Basic Setup

At each date $n = 0, 1, 2, \dots$ there is a state vector $y_n \in \mathbb{R}^{n_y}$, and an agent controls a vector of actions $a_n \in \mathbb{R}^{n_a}$. There are stochastic shocks to both the state evolution and the agent's actions, which follow the linear state space model:

$$y_{n+1} = Ay_n + Ba_n + \Sigma_y W_{n+1} \quad (1)$$
$$a_n = u_n + \Sigma_a W_{n+1}. \quad (2)$$

Here $u_n$ is the part of the actions $a_n$ which is controllable by the agent, $(A, B, \Sigma_y, \Sigma_a)$ are coefficient matrices and $W_{n+1} \in \mathbb{R}^{n_w}$ is an i.i.d. shock vector with distribution $F$. We mostly focus on the case where $F$ is Gaussian, but some of our results are stronger for bounded shocks. Note that we allow for correlation between the shocks to the state in (1) and the actions in (2).

The agent knows that his choices affect the state with some noise as in (2), but he does not know the state evolution (1). Instead he chooses actions based on a subjective model which he updates over time. We allow for the possibility that the agent's model is misspecified, in that it may omit some relevant variables. Thus instead of conditioning on the full vector $y_n$ he may only consider a sub-vector $s_n = K_s y_n$, where $K_s$ selects the appropriate elements. We collect these states and the agent's actions into the agent's state $x_n = [s_n', a_n']' \in \mathbb{R}^{n_x}$. The subjective model is:

$$y_{n+1} = \gamma' x_n + \eta_{n+1}. \quad (3)$$

Here $\gamma \in \mathbb{R}^{n_x \times n_y}$ is a matrix of beliefs or regression coefficients, and $\eta_{n+1}$ is a vector of regression errors. The errors are believed to be orthogonal to the regressors $x_n$:

$$\tilde{E}\left[x_n(y_{n+1} - \gamma' x_n)'\right] = 0. \quad (4)$$

Here $\tilde{E}$ represents the agent's subjective expectation, which may not agree with the objective expectation, particularly if the agent's model is misspecified.[4]

At every date, the agent chooses his actions to maximize his utility given his current beliefs. This is an example of what Kreps (1998) and Sargent (1999) call an "anticipated utility" model: each period the agent makes decisions treating his beliefs $\gamma$ as constant, and then updates the beliefs upon observing outcomes. We suppose the agent has quadratic preferences over his state vector $x_n$, with the positive definite weighting matrix $\Omega(\gamma)$ which may depend on the beliefs, and discount factor $\beta \in (0, 1)$. Thus the agent's problem is:

$$\max_{\{u_n\}} -\frac{1}{2}\tilde{E}\sum_{n=0}^{\infty} \beta^n x_n' \Omega(\gamma) x_n, \quad (5)$$

[4]In practice, the state vector $y_n$ may include variables, such as lagged variables or a constant term, whose evolution would be trivial to learn. Thus the agent need not learn all parts of (1), and thus some elements of $\eta_{n+1}$ may be identically zero.

subject to (2) and (3). The solution is a linear decision rule:

$$u_n = h(\gamma)s_n, \tag{6}$$

where we emphasize the dependence on the beliefs $\gamma$. Substituting (6) into (1) and using (2), we obtain the belief-dependent linear law of motion:

$$y_{n+1} \;=\; [A + Bh(\gamma)K_s]y_n + (\Sigma_y + B\Sigma_a)W_{n+1}. \tag{7}$$

### 2.2. Self-Confirming Equilibrium

Following Fudenberg and Levine (1998) and Sargent (1999), we now define a self-confirming equilibrium as a matrix of beliefs which are consistent with the agent's observations. First, we introduce a bit of notation to summarize the variables entering the orthogonality condition (4). Let $\xi_{n+1} = [y_n', W_{n+1}']'$ and note from (7) that we can write its evolution as:

$$\xi_{n+1} = \begin{bmatrix} A + Bh(\gamma)K_s & \Sigma_y + B\Sigma_a \\ 0 & 0 \end{bmatrix} \xi_n + \begin{bmatrix} 0 \\ I \end{bmatrix} W_{n+1} \equiv \Theta(\gamma)\xi_n + \Lambda W_{n+1}. \tag{8}$$

Then define $g$ as the function whose expectation is zero in (4):

$$g(\gamma, \xi_{n+1}) = x_n(y_{n+1} - \gamma' x_n)' \tag{9}$$

Here we recall that $x_n$ is a linear function of $\xi_{n+1}$ under (6), as is $y_{n+1}$ from (7). Thus $g$ is a quadratic function of $\xi_{n+1}$, a particular structure we will exploit in our analysis.

The key orthogonality condition (4) can then be written as $\tilde{E}g(\gamma, \xi_{n+1}) = 0$. In a self-confirming equilibrium this orthogonality condition holds under the objective probability measure induced by (7) as well. That is, the agent's beliefs are confirmed by his observations. In order for the objective expectation to make sense, we assume that given $\gamma$, $y_n$ has a stationary distribution denoted $\pi$. We later constrain the evolution of beliefs to insure that $\pi$ exists. Thus define $\bar{g}$ as the unconditional expectation of $g$:

$$\bar{g}(\gamma) = E[g(\gamma, \xi_{n+1})] = \int\int g(\gamma, y, W)d\pi(y)dF(W). \tag{10}$$

Definition 2.1. A *self-confirming equilibrium* (SCE) is a matrix $\bar{\gamma} \in \mathbb{R}^{n_x \times n_y}$ such that $\bar{g}(\bar{\gamma}) = 0$.

### 2.3. Adaptation

As we noted above, the agent treats the parameters of his model as constant when making decisions, but then updates them with observations. We specify that the agent learns via the following constant gain recursive least squares algorithm:

$$\gamma_{n+1} \;=\; \gamma_n + \varepsilon R_n^{-1} g(\gamma_n, \xi_{n+1}) \tag{11}$$

$$R_{n+1} \;=\; R_n + \varepsilon\phi\left(x_n x_n' - R_n\right). \tag{12}$$

Here the scalar $\varepsilon > 0$ is the gain, giving the weight on new information relative to the past. The new information is summarized by $g$, whose expectation $\bar{g}$ is zero in a SCE. Thus the algorithm adjusts beliefs in a direction that makes $\bar{g}$ tend toward zero. The new term $R_n$ is an estimate of the second moments of the regressors. More volatile regressors convey less information, and so are given less weight. We introduce the factor $\phi$ in (12) to also include the class of generalized stochastic gradient learning rules, as studied by Evans, Honkapohja, and Williams (2008). These rules set $\phi = 0$ and thus use a constant weighting matrix $R_n = R$.

If the gain decreased over time as $\varepsilon = \frac{1}{n+1}$, then (11)-(12) would be a recursive representation of the standard OLS estimator, as is widely used in the learning literature. A constant gain discounts past observations, implying that the agent pays more attention to more recent data. Such algorithms are known to work well in non-stationary environments, and to provide good predictors even when the underlying model is misspecified.[5] Both motivations are appropriate here, as the agent's model is potentially misspecified and the environment effectively changes as he learns over time.

As noted above, the analysis in this paper presumes that $y_n$ in (7) is stationary. If the agent were to know the true model (1), then it would be sufficient to make the typical stabilizability assumption that there exists a control rule which stabilizes the system. However since the agent's model (3) differs from the truth, and since his beliefs evolve over time, we need to constrain the learning rule so that it does not induce instability in the state evolution. As we cannot guarantee stability in general, we need to confine the estimate sequence to a feasible set. The following assumption restricts beliefs $\gamma$ to guarantee the existence of a unique, stable invariant distribution $\pi$ for $y_n$.

Assumption 2.1. *Let $\mathcal{G} \subset \mathbb{R}^{n_x \times n_y}$ be the set of $\gamma$ such that the eigenvalues of $A + Bh(\gamma)K_s$ have modulus strictly less than one. For each $n$, we assume $\gamma_n \in \mathcal{G}$.*

One way to insure this stability in practice is to impose a projection facility on (11), as in Marcet and Sargent (1989) which restricts the updating rule so that the estimates stay in the set. We will not explicitly deal with such a facility here, as we assume that the SCE $\bar{\gamma}$ is in interior of $\mathcal{G}$ and we analyze escapes to points that remain in the interior of $\mathcal{G}$.[6]

## 3. OVERVIEW OF THE ANALYSIS

This section provides a brief overview of our results. We first show that on average, the agent will be drawn toward a self-confirming equilibrium. Then we characterize events in which beliefs escape from the SCE, which is the main focus of the paper.

[5]Sargent and Williams (2005) discuss the performance of the constant gain algorithm for drifting coefficients. Evans, Honkapohja, and Williams (2009) show that constant gain rules are robust to misspecification.

[6]When $y_n$ contains a constant term, then only $n_y - 1$ eigenvalues of the state matrix in (7) need be less then one.

### 3.1. Convergence of Beliefs

All of our results consider small gain limits. In any sample from the model the gain is constant, thus we look across different samples indexed by the gain. We emphasize this by writing $\gamma_n^\varepsilon$. As $\varepsilon \to 0$ the agent averages more evenly over past data, and the changes in beliefs become smoother. To see this, define the random variable $v_{n+1}^\varepsilon$:

$$v_{n+1}^\varepsilon = (R_n^\varepsilon)^{-1} \left[ g\left(\gamma_n^\varepsilon, \xi_{n+1}\right) - \bar{g}\left(\gamma_n^\varepsilon\right) \right].$$

Then we can re-write (11) as:

$$\frac{\gamma_{n+1}^\varepsilon - \gamma_n^\varepsilon}{\varepsilon} = (R_n^\varepsilon)^{-1} \bar{g}\left(\gamma_n^\varepsilon\right) + v_{n+1}^\varepsilon. \tag{13}$$

Note that (13) is similar to a finite-difference approximation of a time derivative, on a time scale where $\varepsilon$ is the increment between observations. Letting $\varepsilon \to 0$, this approximation becomes arbitrarily good. Along this same limit, a law of large numbers insures that $v_n^\varepsilon$ converges to zero. Thus in the limit we obtain the differential equations:

$$\dot{\gamma} = R^{-1} \bar{g}(\gamma) \tag{14}$$

$$\dot{R} = \phi[\bar{M}(\gamma) - R] \tag{15}$$

Equation (15) carries out a similar limit for (12), where we use the notation $\bar{M}(\gamma) = E(x_n x_n')$. We call these ODEs the *mean dynamics*, as they govern the expected evolution of the agent's beliefs. Theorem 5.1 below makes this formal. Note that an equilibrium point $\bar{\gamma}$ of (14) is a self-confirming equilibrium, and let $M(\bar{\gamma}) = \bar{R}$. Thus if the SCE is stable under the ODE, we see that as $\varepsilon \to 0$ the agent's beliefs (11)-(12) converge to $(\bar{\gamma}, \bar{R})$.

Most of the learning literature has focused on convergence toward an equilibrium, and so has characterized the mean dynamics. Our contribution here is to study rare deviations from an equilibrium, and so we characterize the escape dynamics.

### 3.2. Escape Dynamics

The convergence results show that escapes from the SCE are unlikely. We now show that if they do happen, they are very likely to happen in a particular most probable way and with a predictable frequency. First we define the key objects of interest. To simplify the presentation, we initialize all paths at the SCE.[7] In the following we use the same time scale as (13) where $\varepsilon$ is the time increment between $n$ and $n+1$.

Definition 3.1. Fix an $\varepsilon > 0$, a time horizon $\bar{n} < \infty$ (which may depend on $\varepsilon$), and a compact set $G \subset \mathcal{G}$ with non-empty interior and $\bar{\gamma} \in G$. Let $\gamma^\varepsilon(t), t \in [0, \bar{n}\varepsilon]$ be the piecewise linear interpolation of $\{\gamma_n^\varepsilon\}$.

1. An *escape path* from $G$ is a sequence $\{\gamma_n^\varepsilon, R_n^\varepsilon\}_{n=0}^{\bar{n}}$ solving (11)-(12) such that $\gamma_0^\varepsilon = \bar{\gamma}$ and $\gamma_m^\varepsilon \notin G$ for some $m \leq \bar{n}$. Let $\Gamma^\varepsilon(G, \bar{n})$ be the set of escape paths.

[7] These results can be easily extended to allow for initialization in a neighborhood of the SCE.

2. For any sequence $\{\gamma_n^\varepsilon, R_n^\varepsilon\}_{n=0}^{\bar{n}}$ solving (11)-(12) with $\gamma_0 = \bar{\gamma}$, define the (first) *escape time* from $G$ as:

$$\tau^\varepsilon(\{\gamma_n^\varepsilon\}) = \varepsilon \inf\{m : \gamma_m^\varepsilon \notin G\} \in \mathbb{R} \cup \{\infty\}.$$

3. A *regular escape path* from $G$ is an escape path for which $\exists \mu_2 > \mu_1 > 0$ with $\{\gamma : \|\gamma - \bar{\gamma}\| < \mu_2\} \in G$ such that there is no $t'' > t'$ where $\|\gamma^\varepsilon(t') - \bar{\gamma}\| > \mu_2$ and $\|\gamma^\varepsilon(t'') - \bar{\gamma}\| \leq \mu_1$. Let $\bar{\Gamma}^\varepsilon(G, \bar{n})$ be the set of regular escape paths.

For small gains, any path $\{\gamma_n^\varepsilon\}$ spends most of its time near the SCE $\bar{\gamma}$, and if noise pushes it away, it tends to be drawn back. While eventually all paths leave the set $G$, an escape path exits before the terminal date $\bar{n}$. A regular escape path is one which upon escaping from a $\mu_2$ neighborhood of the SCE does not return to a smaller $\mu_1$ neighborhood of it. Once a regular escape path starts to escape, it does not turn back.

Our results characterize bounds on the probability of escape, the mean escape time, and the *most probable escape path*.[8] In particular, we show below that almost all escape paths exit the set $G$ at the end of the most probable escape path, and almost all regular escape paths are close to the most probable path while they are still in the set $G$.

We characterize escapes as arising due to a perturbation $v$ of the mean dynamics:

$$\dot{\gamma} = R^{-1}\bar{g}(\gamma) + v. \tag{16}$$

Since we focus on the beliefs $\gamma$, along these escape paths we let $R$ follow the mean dynamics (15) conditional on $\gamma$. Note the similarity of (16) to (13): for the mean dynamics the perturbation vanishes, but it resurfaces to govern the escape dynamics. That is, the perturbation $v_n^\varepsilon$ in (13) has a zero mean and so it typically becomes negligible and the beliefs track the mean dynamics. But to characterize the unlikely sequence of shocks leading to an escape, we analyze the perturbations $v$ in (16) which cause the beliefs to escape. Alternative escape paths are associated with alternative perturbations, and we evaluate the likelihood of alternative escape paths by a cost function which penalizes less likely perturbations.

The "cost" of a particular perturbation depends on its size relative to the volatility of beliefs, as bigger perturbations will more naturally occur with more volatility. If $y_n$ (and hence $\xi_n$) were i.i.d., then for a given $(\gamma, R)$ the variance of the belief increment in (13) would measure belief volatility:

$$Var(v_{n+1}) = R^{-1}E\left[g(\gamma, \xi_{n+1})\, g(\gamma, \xi_{n+1})' - \bar{g}(\gamma)\, \bar{g}(\gamma)'\right] R^{-1}. \tag{17}$$

However since $y_n$ is temporally dependent, we need a more general measure. Escapes may occur either due to a few large shocks or a succession of small shocks, and the dependence of $y_n$ is crucial in evaluating the likelihood of such alternative perturbations. As discussed in Meyn and Tweedie (1993, Ch. 17), one way to capture this

[8]Our definition of a regular escape path follows Freidlin and Wentzell (1998) and Dupuis and Kushner (1987). Our notion of the most probable escape path follows Maier and Stein (1997). Cho, Williams, and Sargent (2002) call this a dominant escape path.

dependence is to first find the solution $f$ of the Poisson equation:

$$f(\gamma,\xi) - E[f(\gamma,\xi_{n+1})|\xi_n = \xi] = g(\gamma,\xi) - \bar{g}(\gamma). \tag{18}$$

Then for fixed $\gamma$ the left side of (18) is a martingale, which simplifies the characterization of the long run belief variance. In particular, for fixed $(R,\gamma)$ we have:

$$\frac{1}{\sqrt{T}}\sum_{n=0}^{T-1} R^{-1}[g(\gamma,\xi_n) - \bar{g}(\gamma)] = \frac{1}{\sqrt{T}}\sum_{n=0}^{T-1} R^{-1}\left(f(\gamma,\xi_n) - E[f(\gamma,\xi_{n+1})|\xi_n]\right)$$

Using martingale convergence results, Meyn and Tweedie (1993) show that as $T \to \infty$ these sums converge to a normal random vector with covariance matrix:

$$Q(\gamma,R) = R^{-1}E\left(f(\gamma,\xi)f(\gamma,\xi)' - E\left[f(\gamma,\xi_{n+1})|\xi_n = \xi\right]E\left[f(\gamma,\xi_{n+1})|\xi_n = \xi\right]'\right)R^{-1}. \tag{19}$$

This same matrix $Q$ is our measure of the belief volatility. When $\xi_n$ is i.i.d., then $f = g$ and $Q$ is equal to $Var(v_n)$ as in (17). But in general, we need to solve the Poisson equation (18). While this is not always an easy task, Lemma 6.1 in Section 6.2 below shows how to explicitly solve (18) for our model.

We analyze escapes on a fixed continuous time horizon $\bar{T} < \infty$, and set $\bar{n} = \bar{T}/\varepsilon$. Thus $\bar{n} \to \infty$ as $\varepsilon \to 0$. Also $Q$ may be singular, so let $Q^{\dagger}$ be its pseudoinverse. To characterize the escape dynamics, we choose the perturbations $v$ in (16) which push beliefs to the boundary $\partial G$ in the most cost effective way:

$$\bar{S} = \inf_{v(\cdot),T}\frac{1}{2}\int_0^T \sum_{i=1}^{n_y} v_i(s)'Q(\gamma(s),R(s))^{\dagger}v_i(s)ds \tag{20}$$

where $v_i$ is a column of $v$ and the minimization is subject to (16), (15), and:

$$\gamma(0) \ = \ \bar{\gamma}, \ R(0) = \bar{R}, \ \gamma(T) \in \partial G \text{ for some } 0 < T \leq \bar{T}. \tag{21}$$

Thus the instantaneous cost of a perturbation is quadratic, with weighting matrix $Q$. If $v \equiv 0$ then the beliefs follow the mean dynamics. The cost is zero, but the beliefs do not escape. To find the most probable escape path, we find a least cost path of perturbations that pushes beliefs from $\bar{\gamma}$ out to the boundary $G$.

The characterization of the escape dynamics by the control problem (20) is the main result in this paper. We build on the results of Dupuis and Kushner (1989), who characterize escape dynamics in models like ours as a general variational problem. But their results are difficult to directly apply. In this paper, we derive the explicit form of the control problem in (20) for our class of models, which makes the theory directly applicable. We discuss the details of our characterization in Section 6 below.

The following theorem shows how (20) characterizes escapes. We fix a set $G$ and horizon $\bar{T}$ as above, and recall that $\tau^{\varepsilon}$ is an escape time, with the escape taking place at $\gamma^{\varepsilon}(\tau^{\varepsilon})$. We say that the minimized cost $\bar{S}$ is continuous in $G$ if we obtain the same value when we change the terminal condition in (21) to an interior point arbitrarily close to the boundary of $G$.[9] The additional necessary conditions A.1 and A.3 are in Appendix A.1 and A.2.2, respectively.

THEOREM 3.1. *Suppose that Assumptions 2.1, A.1, and A.3 hold, let $\gamma^\varepsilon(\cdot)$ be the piecewise linear interpolation of $\{\gamma_n^\varepsilon\}$, and let $\gamma(\cdot) : [0, \bar{T}] \to \mathbb{R}^{n_x \times n_y}$ solve (20).*

*1. Suppose that the shocks $W_n$ are i.i.d. and unbounded (but have exponential tails). Then we have:*

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

*2. Suppose that the shocks $W_n$ are i.i.d. and bounded, and $\bar{S}$ is continuous in $G$. Then we have:*

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

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

$$\lim_{\varepsilon \to 0} P\left[\exp\left((\bar{S} + \delta)/\varepsilon\right) > \tau^\varepsilon > \exp\left((\bar{S} - \delta)/\varepsilon\right)\right] = 1,$$
$$\text{and: } \lim_{\varepsilon \to 0} \varepsilon \log E(\tau^\varepsilon) = \bar{S}.$$

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

$$\lim_{\varepsilon \to 0} P\left(\|\gamma^\varepsilon(\tau^\varepsilon) - \gamma(T)\| < \delta | \{\gamma_n^\varepsilon\} \in \Gamma^\varepsilon(G, \bar{n})\right) = 1.$$
$$\text{Moreover: } \lim_{\varepsilon \to 0} P\left(\|\gamma^\varepsilon(t) - \gamma(t)\| < \delta,\ t < \tau^\varepsilon(\{\gamma_n^\varepsilon\}) | \{\gamma_n^\varepsilon\} \in \bar{\Gamma}^\varepsilon(G, \bar{n})\right) = 1.$$

*Proof.* See Appendix A.2.2. ∎

Part (1) shows that the probability of observing an escape on a bounded time interval is exponentially decreasing in the gain $\varepsilon$, with the rate given by the minimized cost function $\bar{S}$. The next three parts establish stronger results for bounded shocks.[10] Part (2) shows that in this case the asymptotic inequality in part (1) becomes an equality. Part (3) shows that for small $\varepsilon$ the escape times from the SCE become close to $\exp(\bar{S}/\varepsilon)$.[11] The log mean escape time also converges to this value. Finally, part (4) shows that the minimizing path from (20) is the most probable escape path. This means that with probability approaching one, all escapes occur near the end of this path, and all regular escape paths remain near it.

[9] More precisely, let $\bar{S}_\delta$ be the value obtained in (20) when we change (21) to require that $\|\gamma(T) - \gamma^*\| < \delta$ for some $\gamma^* \in \partial G$. Then $\bar{S}$ is continuous in $G$ if $\lim_{\delta \to 0} \bar{S}_\delta = \bar{S}$.

[10] Although we focus mainly on Gaussian shocks, our results are sharpest in the bounded case. An application of these results in Cho, Williams, and Sargent (2002) obtained nearly identical results with bounded and unbounded shocks. Thus many results may carry over for some unbounded cases.

[11] But notice that $\delta$ is fixed, and thus as $\varepsilon \to 0$ the interval around $\exp(\bar{S}/\varepsilon)$ expands.

## 4. AN EXAMPLE

We now present a simple example which shows how escape dynamics may be a dominant feature of a model. The example here shares many features with other applications with prominent escape dynamics, as we discuss in Section 4.4 below. The model has a unique self-confirming equilibrium which is stable under learning. In one parameterization of the model, simulations from the model show small fluctuations around the SCE, with escapes being the rare large movements away from it. However under a different parameterization, simulations show recurrent escapes from the SCE which always lead in the same direction. We apply our methods to characterize the escapes and to explain these differences.

### 4.1. The Model

We study a model of a monopolist learning its demand curve. The firm faces an unknown linear demand for its products, which is subject to a shock. The firm produces at a constant cost per unit in each period, with the realized cost in a period depending on a shock. In one parameterization we consider the shocks to costs and demand are uncorrelated, while in the other they are correlated. For simplicity, the model is mostly static, with dynamics only coming through learning.

In particular, we specialize the general model of Section 2.1 as follows. The state vector consists of output $d_n$ and a constant: $y_n = [d_n, 1]'$, and the shock vector $W_n = [W_{1n}, W_{2n}]'$ is a 2 dimensional standard normal random vector. For simplicity we set the expected marginal cost to zero, so we can think of the firm choosing its markup $u_n$ over marginal cost, with the cost shock determining the realized price $a_n$:

$$a_n = u_n + \sigma_a W_{2,n+1}.$$

Output is given by the static linear demand curve:

$$d_{n+1} = b_0 + b_1 a_n + \sigma_y W_{1,n+1} + \rho\sigma_a W_{2,n+1}.$$

Note the dating convention: $d_{n+1}$ is the current period's output, which depends on the current price $a_n$. When $\rho \neq 0$, the shocks to costs and demand are correlated. Thus output and prices are determined by (1) and (2) with the following specification:

$$A = \begin{bmatrix} 0 & b_0 \\ 0 & 1 \end{bmatrix},\ B = \begin{bmatrix} b_1 \\ 0 \end{bmatrix},\ \Sigma_y = \begin{bmatrix} \sigma_y & \rho\sigma_a \\ 0 & 0 \end{bmatrix},\ \Sigma_a = [0, \sigma_a].$$

The firm does not know its demand curve, but instead sets its price based on its subjective model (3), which here takes the form:

$$d_{n+1} = \gamma_0 + \gamma_1 a_n + \eta_{n+1} \tag{22}$$

thus $s_n = [0, 1]y_n = 1$ and $x_n = [1, a_n]'$. We abstract from the second equation of the state evolution (1), as it is simply an identity with nothing to learn. Note that when $\rho \neq 0$, the belief equation (22) is a misspecified regression for (1), as $a_n$ and $d_{n+1}$

are driven by correlated shocks. The firm maximizes expected profits based on (22), which can be written as in (5):

$$\tilde{E}_n[a_n d_{n+1}] = \tilde{E}_n[\gamma_0 a_n + \gamma_1 a_n^2 + \eta_{n+1} a_n] = \tilde{E}_n x_n' \begin{bmatrix} 0 & \gamma_0/2 \\ \gamma_0/2 & \gamma_1 \end{bmatrix} x_n,$$

where we use (4), and the last equation implicitly defines the weighting matrix $\Omega(\gamma)$. Since there are no dynamics in the model, the optimization problem (5) is static. Thus the policy function (6) determines the optimal markup:

$$u_n = h(\gamma) \equiv -\frac{\gamma_0}{2\gamma_1}.$$

### 4.2. Learning

For simplicity, we study here a simple generalized stochastic gradient algorithm which sets $\phi = 0$ in (12) and $R_n \equiv \bar{R}$ in (11). As discussed above, Evans, Honkapohja, and Williams (2009) study such rules and relate them to the recursive least squares rule. The key function $g = [g_1, g_2]'$ from (9) is given explicitly here by:

$$\begin{aligned} g_1(\gamma,\xi) &= b_0 - \gamma_0 - (b_1 - \gamma_1)\frac{\gamma_0}{2\gamma_1} + (b_1 + \rho - \gamma_1)\sigma_a W_2 + \sigma_y W_1 \\ g_2(\gamma,\xi) &= g_1(\gamma,\xi)\left(-\frac{\gamma_0}{2\gamma_1}\right) + \left(b_0 - \gamma_0 - (b_1 - \gamma_1)\frac{\gamma_0}{2\gamma_1}\right)\sigma_a W_2 + (b_1 + \rho - \gamma_1)\sigma_a^2 W_2^2 + \sigma_a\sigma_y W_1 W_2. \end{aligned}$$

Since $\xi = [1, W']'$ is i.i.d., we simply take expectations to get $\bar{g} = [\bar{g}_1, \bar{g}_2]'$ as in (10):

$$\begin{aligned} \bar{g}_1(\gamma) &= b_0 - \gamma_0 - (b_1 - \gamma_1)\frac{\gamma_0}{2\gamma_1} \\ \bar{g}_2(\gamma) &= \bar{g}_1(\gamma)\left(-\frac{\gamma_0}{2\gamma_1}\right) + (b_1 + \rho - \gamma_1)\sigma_a^2 \end{aligned}$$

Since the model is static, Assumption 2.1 which insures stability is satisfied.[12] It is straightforward to show that there is a unique self-confirming equilibrium, given by:

$$\bar{\gamma} = [\bar{\gamma}_0, \bar{\gamma}_1]' = \left[\frac{2b_0(b_1+\rho)}{2b_1+\rho}, b_1 + \rho\right]'.$$

Note that when $\rho \neq 0$ the SCE beliefs are biased estimates of the true intercept and slope $(b_0, b_1)$. By our discussion above, we find that as $\varepsilon \to 0$ the beliefs converge to the SCE $\bar{\gamma}$.[13] Thus we expect that for small $\varepsilon$, beliefs will remain near the SCE $\bar{\gamma}$, and the price will exhibit small fluctuations around the expected price $h(\bar{\gamma})$.

[12]Some of our results are simplified if prices are guaranteed to be positive. In practice we require the firm's estimated demand curve to slope downward, that is $\gamma_0 \geq 0$ and $\gamma_1 < K < 0$ for some small $K$. Such a constraint was never binding in our calculations or simulations.

[13]This analysis is made formal in Theorem 5.1 below. Appendix A.3 provides formal detail and verifies all necessary assumptions for the example.

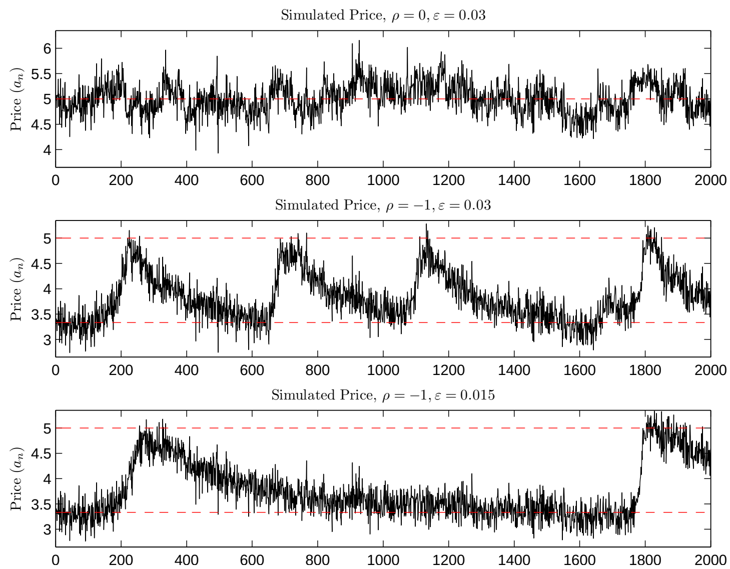


**FIGURE 1.** Simulated price series $a_n$ for two parameterizations of the model: $\rho = 0$ (top panel) and $\rho = -1$ (bottom panel). The red dashed lines show the SCE expected prices when $\rho = 0$ ($a_n = 5$) and when $\rho = -1$ ($a_n = 3\frac{1}{3}$).

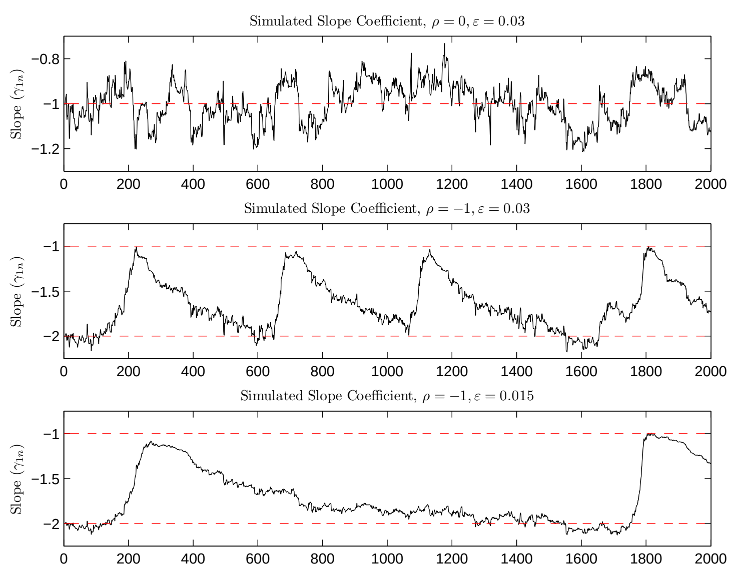


**FIGURE 2.** Simulated slope coefficient series $\gamma_{1n}$ for two parameterizations of the model: $\rho = 0$ (top panel) and $\rho = -1$ (bottom panel). The red dashed lines show the SCE slopes when $\rho = 0$ ($\bar{\gamma}_1 = b_1 = -1$) and when $\rho = -1$ ($\bar{\gamma}_1 = b_1 + \rho = -2$).

However we find that the model exhibits some striking behavior, as shown in Figures 1 and 2, which plot some simulated outcomes from the model for different settings of $\rho$. Figure 1 plots simulated time paths of prices $a_n$ for $\rho = 0$ and $\rho = -1$, while Figure 2 plots the estimated slope coefficients $\gamma_{1n}$ from the same simulations.[14] Each figure also plots two different gain settings when $\rho = -1$. When $\rho = 0$, and thus the regression model (22) is correctly specified, the price and belief series behave as expected. The price is near $h(\bar{\gamma})$ and the slope is near $\bar{\gamma}_1$ throughout, with movements away from these levels following no regular pattern. But when $\rho = -1$, we observe recurrent episodes in which the firm raises its price sharply, corresponding to the episodes in which the slope coefficient increases from its SCE level of $b_1 + \rho = -2$ to a value near the true slope of $b_1 = -1$. The higher price and lower slope are sustained for a relatively short time, as they gradually are drawn back to the SCE levels. As our results suggest, for smaller gain settings the escapes become less frequent and the model spends an increasing fraction of time near the SCE. But escapes do recur, and the escape paths have a very regular pattern, leading to increases in the slope and the price. We now show that this is precisely what our results predict.

### 4.3. Escape Dynamics

To analyze the escapes, we first calculate the cost function matrix $Q$ from (19). Since the model is static and the shocks are i.i.d., the Poisson equation (18) is trivial and $f = g$. Thus $Q$ is the covariance matrix in (17), which can be written here as:

$$Q(\gamma) = \bar{R}^{-1} \begin{bmatrix} Q_{11}(\gamma) & Q_{12}(\gamma) \\ Q_{12}(\gamma) & Q_{22}(\gamma) \end{bmatrix} \bar{R}^{-1},$$

where: $Q_{ij}(\gamma) = E\left[g_i(\gamma,\xi)g_j(\gamma,\xi) - \bar{g}_i(\gamma)\bar{g}_j(\gamma)\right]$. Using the properties of normal random variables we have explicitly:

$$\begin{aligned} Q_{11}(\gamma) &= (b_1 + \rho - \gamma_1)^2\sigma_a^2 + \sigma_y^2 \\ Q_{12}(\gamma) &= (b_1 + \rho - \gamma_1)\left(\bar{g}_1(\gamma) + h(\gamma)(b_1 + \rho - \gamma_1)\right)\sigma_a^2 + h(\gamma)\sigma_y^2 \\ Q_{22}(\gamma) &= (\bar{g}_1(\gamma) + h(\gamma)(b_1 + \rho - \gamma_1))^2\sigma_a^2 + h(\gamma)^2\sigma_y^2 + 2(b_1 + \rho - \gamma_1)^2\sigma_a^4 + \sigma_a^2\sigma_y^2 \end{aligned}$$

We now solve the control problem (20) to characterize the escapes. Rather than fixing a single set $G$, we consider sets of the form:

$$G(r) = \left\{ \gamma : \|\gamma - \bar{\gamma}\| < |r|, \quad \frac{r}{|r|}\gamma_1 > \frac{r}{|r|}\bar{\gamma}_1 \right\}.$$

That is with $r > 0$, $G(r)$ is the half-ball of radius $r$ around $\bar{\gamma}$ with $\gamma_1 > \bar{\gamma}_1$, while $G(-r)$ is the other half-ball with $\gamma_1 < \bar{\gamma}_1$. The minimized cost $\bar{S}$ for different radii is shown in Figure 3 when $\rho = 0$ and when $\rho = -1$. When $\rho = 0$ we see that $\bar{S}$ is quite symmetric, increasing rapidly once we move away from zero in either direction. Thus escapes of substantial size become quite rare for small gains, and the escapes that do occur are

[14] The other parameters in the model are set as follows: $b_0 = 10,\ b_1 = -1, \sigma_a = \sigma_y = 0.2$.

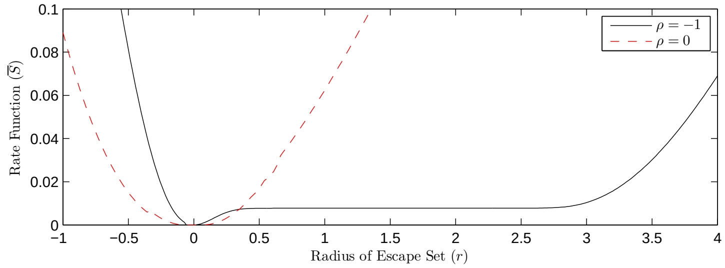


**FIGURE 3.** Rate function $\bar{S}$ for different escape radii $r$ under two parameterizations of the model: $\rho = -1$ (black solid line) and $\rho = 0$ (red dashed line).

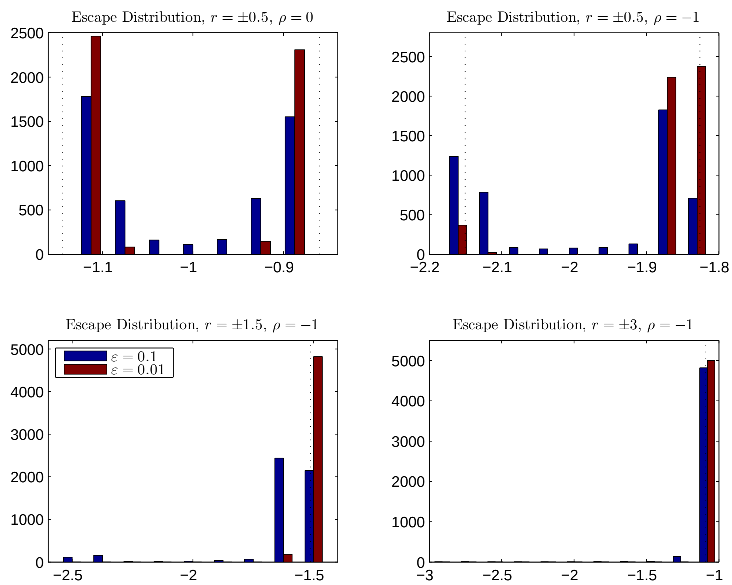


**FIGURE 4.** Escape distributions for the slope coefficient $(\gamma_1^\varepsilon(\tau^\varepsilon))$ for different gains $(\varepsilon)$, different radii $(r)$, and two settings of $\rho$. Each panel considers a different $(\rho, r)$ combination and plots the distribution for $\varepsilon = 0.1$ (blue bars) and $\varepsilon = 0.01$ (red bars). The black dotted lines show the terminal point of the most probable escape $(\gamma_1(T))$.

equally likely to result in increases or decreases of the slope coefficient. However when $\rho = -1$ the minimized cost function is asymmetric. It increases rapidly in the negative direction, but it has a long, nearly flat section associated with increases in the slope coefficient $(r > 0)$. Thus even though escapes are rare, once beliefs do escape they are likely to move quite a distance (nearly three Euclidean units) in the positive direction. This reflects what we observed in Figure 2, where with $\rho = 0$ the movements away from the SCE followed no definite pattern, while with $\rho = -1$ there were recurrent escapes with increases in the slope coefficient of roughly the same magnitude.

A more comprehensive look at escapes is given in Figure 4, which plots escape distributions for different $r$ and $\rho$ settings. In particular, each panel of the figure fixes $(r, \rho)$ and shows a histogram of the slope coefficient at the time of escape $(\gamma_1^\varepsilon(\tau^\varepsilon))$ from 5000 simulated escapes for each of two different gain settings, $\varepsilon = 0.1$ and $\varepsilon = 0.01$.[15]

[15] In each simulation for this and the following figures, we initialize beliefs at the SCE, and stop once the beliefs exit the set $G(r)$ or the time reaches $\bar{n} = 500/\varepsilon$.

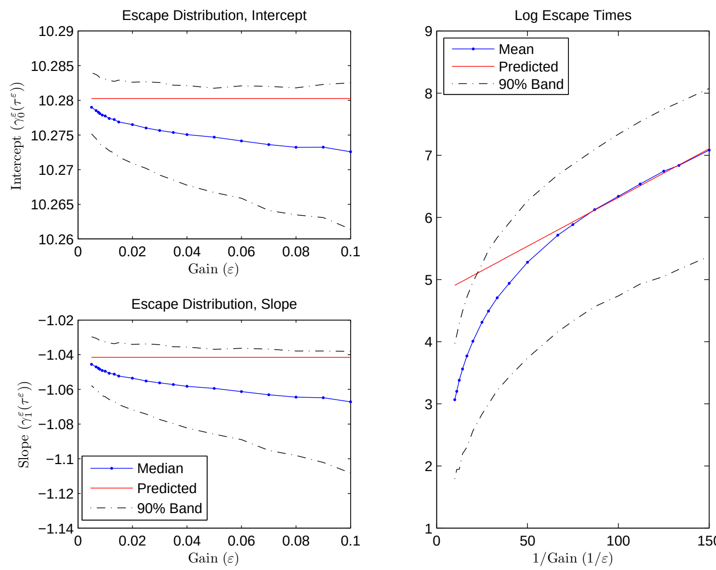


**FIGURE 5.** Escape distributions and escape times. The left panels plot the escape distribution of the coefficient $(\gamma_0^\varepsilon(\tau^\varepsilon), \gamma_1^\varepsilon(\tau^\varepsilon))$ for different gains ($\varepsilon$), fixing $r = 3.2$ and $\rho = -1$. Each panel plots the median (blue dot-lines) and 90% band (black dashed lines) of the simulated distribution along with the predicted escape point (red solid lines). The right panel plots the log escape times, showing the log mean time (blue dot-lines) and 90% band (black dashed lines) of the simulated distribution, along with the predicted log times (red solid line).

Also shown in each figure is the predicted escape point $\gamma_1(T)$ from the most probable escape path associated with that particular $(r, \rho)$ setting. The top two panels consider relatively small escapes, with $r = \pm 0.5$. There we see that the escape distributions are relatively symmetric, with escapes in both the positive and negative directions. The middle of the distribution thins out as the gain decreases, with the escapes becoming more concentrated near the predicted points. When $\rho = 0$ the escapes are quite symmetric, but for $\rho = -1$ positive movements are somewhat more likely, becoming more so for smaller gain. The asymmetries that we observed in Figure 3 become even more pronounced when we consider larger escape sets of $r = \pm 1.5$ and $r = \pm 3$ in the bottom panels of Figure 4. Here we only consider $\rho = -1$, as in the simulations we never observed escapes of this magnitude for $\rho = 0$. The panels clearly show that for small gain the escape distributions become highly concentrated around the predicted escape point in the positive direction. Thus Figure 4 illustrates that the asymmetries in the rate functions are borne out in the simulations, and the escape distributions concentrate near the predictions, just as Theorem 3.1 suggests.

We now investigate the predictions of Theorem 3.1 in more detail. Figure 5 plots the escape distributions and escape times when $\rho = -1$ and $r = 3.2$ for varying gains, with 5000 simulations for each gain setting. Part (4) of the theorem says that as $\varepsilon \to 0$ the distribution of escape points should concentrate at the end point of the most probable path. This is documented in the left panels of Figure 5, where we plot the predicted escape point for both the intercept coefficient $\gamma_0$ (top panel) and slope coefficient $\gamma_1$ (bottom panel), along with the median and bands covering 90%

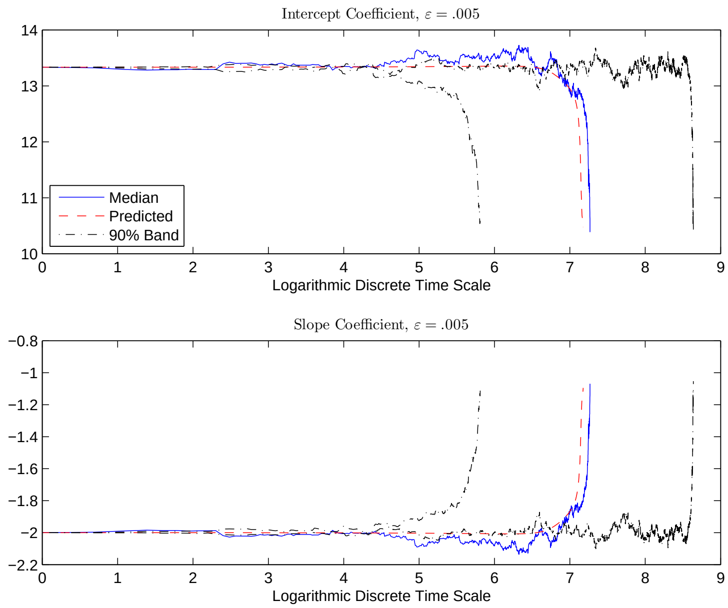


**FIGURE 6.** Simulated and predicted escape paths with $r = 3$, $\rho = -1$ and $\varepsilon = 0.05$ for the intercept coefficient $\gamma_0$ (top panel) and slope coefficient $\gamma_1$ (bottom panel). Each panel plots the median (blue solid lines) and 90% band (black dashed lines) of the simulated distribution along with the predicted escape path (red dashed lines).

of the simulated distribution. Here we clearly see that as $\varepsilon \to 0$ the distribution tightens substantially around the prediction, and the median converges up toward the predicted level. In addition, part (3) of Theorem 3.1 predicts that for small gain the mean escape times increase exponentially in $1/\varepsilon$ at rate $\bar{S}$. The right panel of Figure 5 plots the log mean escape times from the simulations along with our prediction, and bands covering 90% of the simulated distribution. Note that we only predict the slope of the line shown, which gives the exponential rate of increase.[16] For relatively large gains, the escapes occur more rapidly than our results suggest, but for small gains our results provide a good match to what we observed in the simulations.

Finally, part (4) of Theorem 3.1 states that in the limit all regular escapes remain close to the most probable path. This is illustrated in Figure 6, which shows a summary of 5000 simulated escape paths with $r = 3$, $\rho = -1$, and $\varepsilon = 0.005$. The top panel plots the time paths of the intercept coefficient $\gamma_0$ along an escape, while the bottom panel plots the slope. In each case, we plot most probable path resulting from our calculation (20), along with the paths corresponding to the 5%, 50% (median), and 95% quantiles of the simulated escape time distribution. The plots use a logarithmic discrete time scale, so that the most probable path is scaled as $\log(t/\varepsilon)$. The figure shows that all the escape paths are characterized by a long period near the SCE, followed by a rapid increase in the slope and decrease in the intercept. Moreover, our predicted most probable escape path is quite close to the median from the simulations. Thus our results accurately predict the entire time path of beliefs during an escape.

[16]The theorem states that $\log E\tau^{\varepsilon} \approx S_0 + \bar{S}/\varepsilon$ for some constant $S_0$. In the figure the constant was chosen to give a good fit.

### 4.4. Interpretation and Relation to the Literature

The model exhibits strikingly different behavior depending on the value of $\rho$. When $\rho = 0$ escapes are symmetric, leading equally to increases or decreases in the slope coefficient, and escapes become very rare with small gain. However when $\rho = -1$ escapes are asymmetric, leading to regular increases in the slope coefficient from $\gamma_1 = b + \rho$ to $\gamma_1 = b$, and these escapes recur with small gain (albeit less often). As we noted above, when $\rho \neq 0$ the estimate of the slope coefficient in the SCE is biased, due to the correlation between prices and output. This misspecification plays a key role in generating the prominent escape dynamics we observe. Even though the shocks $W_1$ and $W_2$ are independent processes, there will occasionally be a string of correlated realizations. With $\rho < 0$, during such an episode there will be a smaller response of output to the changes in the price. The firm interprets these outcomes as a decrease in the elasticity of its demand curve, and responds by raising its markup $u_n = h(\gamma)$, and hence the mean price. As the firm increases its markup, it gets more influential observations and hence obtains a better estimate of the true slope of its demand curve, which indeed is more inelastic than the SCE suggests. This process is thus self-reinforcing and leads to an escape, but it takes an unlikely sequence of shock realizations in order for it to start. The escapes end when the firm obtains the correct slope estimate, and it does not increase its markup any further. Eventually, there are sufficient independent realizations of the shocks $W_n$, which allow the firm to once again pick up on the correlation between prices and demand, drawing the firm back to the SCE. As this process repeats over time, it generates the episodes of rapidly rising and gradually falling prices we observed in the simulations.

Setting $\rho = 0$ eliminates the misspecification and hence the mechanism leading to the escapes. A correlated string of shock realizations will still cause the firm to lower its estimated demand elasticity, and hence raise its markup. Again this leads to a better estimate of the true slope, but this now agrees with the SCE, which is in fact more *elastic* than the atypical string of shocks suggested. This counteracts the initial effects of the shocks, and leads the firm back to the SCE. Thus the model with $\rho = 0$ lacks the self-reinforcing dynamics which drive the escapes.

Locally self-reinforcing dynamics have been a crucial feature of models with prominent escape dynamics. Most directly, models with multiple equilibria have locally reinforcing dynamics around each equilibrium. These models tend to experience escapes when shocks push beliefs from one equilibrium to another. Notable examples include evolutionary games, such Kandori, Mailath, and Rob (1993) and Young (1993) and much subsequent literature, and macroeconomic models with multiple stable equilibria, such as Kasa (2004). Williams (2002) uses the methods of this paper to study multiple equilibria in a model of learning in games. Similarly, some models feature one stable equilibrium and an explosive region, such as Marcet and Nicolini (2003) and Sargent, Wiliams, and Zha (2008). Here shocks occasionally force beliefs to enter the explosive region, which generates a self-reinforcing acceleration of beliefs. More closely related to this paper are the learning models with a with a unique (self-confirming) equilibrium, where prominent escape dynamics have gone along with misspecifications.

A notable example is Sargent (1999), which was analyzed by Cho, Williams, and Sargent (2002) using the results of this paper. In that model a government sets monetary policy using a misspecified model which does not account for the role of inflation expectations. Related models include Bullard and Cho (2005), McGough (2006), Ellison and Yates (2007), Sargent, Williams, and Zha (2006), and Cho and Kasa (2008), several of which use the methods developed here. Similarly, Williams (2009) considers a duopoly version of the example here, in which firms do not account for the actions of their competitors, and escapes lead to episodes resembling price wars. In all of these cases, the escape dynamics are driven by occasional sequences of shocks which trigger actions that allow agents to temporarily overcome the misspecification of their models.

## 5. CONVERGENCE ANALYSIS

In the rest of the paper we provide more formal detail to support the results in Section 3. In this section we formalize Section 3.1, providing conditions which insure that agents' beliefs converge to a self-confirming equilibrium. The results in this section follow from Kushner and Yin (1997), and are analogous to results in Evans and Honkapohja (2001), with related results in much of the learning literature.

To proceed with the analysis, we define a time scale to convert the discrete time belief evolution into a continuous time process. We let the $\varepsilon$ be the continuous time interval and interpolate between the discrete iterations in the learning rule (11)-(12):

$$\begin{aligned} \gamma^{\varepsilon}(t) &= \theta^{\varepsilon}_{n},\ t \in [n\varepsilon, (n+1)\,\varepsilon) \\ R^{\varepsilon}(t) &= R^{\varepsilon}_{n},\ t \in [n\varepsilon, (n+1)\,\varepsilon) \end{aligned}$$

This defines the continuous time processes as piecewise constant functions of $t \in [0, +\infty)$, which are right-continuous with a left-limit (RCLL). The results in this section establish the weak convergence of these processes on the Skorohod space $D[0, +\infty)$ of RCLL functions. Note that as $\varepsilon \to 0$ the time interval between observations shrinks, and the process becomes smoother. That is, the constant segments become shorter and there are more observations in any given (continuous) time interval. The next theorem shows that as $\varepsilon \to 0$ the interpolated processes converge to the solution of the ODEs we derived informally in (14)-(15).

In Appendix A.1 we provide a proof of the following theorem, and list its necessary conditions as Assumptions A.1. They consist of regularity conditions on the algorithm and the error distribution, and we show there that many of the conditions are satisfied in our baseline model. The remaining conditions require that $\bar{g}(\gamma)$ and $\bar{M}(\gamma)$ be continuous and that the system of ODEs (14)-(15) have an asymptotically stable point $(\bar{\gamma}, \bar{R})$. As we've seen, verifying these conditions is straightforward in practice.

THEOREM 5.1. *Under Assumptions A.1, as $\varepsilon \to 0$, $(\gamma^\varepsilon(\cdot), R^\varepsilon(\cdot))$ converge weakly to $(\gamma(\cdot), R(\cdot))$, where:*

$$\begin{aligned}\gamma(t) &= \gamma(0) + \int_0^t R(s)^{-1}\bar{g}(\gamma(s))ds, \\ R(t) &= R(0) + \int_0^t \phi[\bar{M}(\gamma(s)) - \gamma(s)]ds.\end{aligned}$$

As the ODEs have a stable point at the SCE, the theorem shows that as $\varepsilon \to 0$ over time ($t \to \infty$) agents' beliefs converge weakly to the SCE. Interestingly, the same limiting ODE characterizes the limit of beliefs with decreasing gain algorithms, such as the usual recursive least squares algorithm studied by Marcet and Sargent (1989) which sets $\varepsilon = \frac{1}{n+1}$ in (11)-(12). But with decreasing gain the beliefs typically converge with probability one as $n \to \infty$, while we obtain weak convergence with constant gain. The weaker notion of convergence here means that for any given gain $\varepsilon$, occasional departures from the SCE may persist over time.

## 6. LARGE DEVIATIONS AND ESCAPE DYNAMICS

In this section, we use techniques from large deviation theory to analyze the escape dynamics. The convergence results above show that any event in which beliefs get far from the SCE must have a probability converging to zero with $\varepsilon$. However, for a fixed $\varepsilon > 0$ we may observe such rare escapes, and large deviation theory allows us to characterize them. In this section we provide the details behind our key Theorem 3.1 above, which characterizes the escape dynamics by a deterministic control problem.

### 6.1. Large Deviations

The theory of large deviations deals with calculating bounds on the asymptotic probabilities of rare events. These events have probabilities that converge to zero exponentially fast, and large deviation results identify the exponential rate of convergence. As such, large deviations can be viewed as refinements of classical laws of large numbers and central limit theorems. We now formally define some terminology. Let a sequence $\{X_T\}$ be defined on a probability space $(\Omega, \mathcal{F}, P)$ and take values in a complete separable metric space $\mathcal{X}$. A *rate function* $S_X : \mathcal{X} \to [0, \infty]$ has the property that for any $M < \infty$ the level set $\{x \in \mathcal{X} : S_X(x) \leq M\}$ is compact. We now formally define a large deviation.

DEFINITION 6.1. A sequence $\{X_T\}$ satisfies a *large deviation principle* on $\mathcal{X}$ with *rate function* $S_X$ if the following two conditions hold.

1. For each closed subset $F$ of $\mathcal{X}$, $X_T$ satisfies the large deviation upper bound:

$$\limsup_{T\to\infty} \frac{1}{T} \log P\{X_T \in F\} \leq -\inf_{x\in F} S_X(x).$$

2. For each open subset $G$ of $\mathcal{X}$, $X_T$ satisfies the large deviation lower bound:

$$\liminf_{T\to\infty} \frac{1}{T} \log P\left\{X_T \in G\right\} \geq -\inf_{x\in G} S_X(x).$$

If the sequence converges to a limit not contained in $F$ and $G$, the probability the sequence enters these sets converges to zero. If the sequence satisfies a large deviation principle, this convergence is exponential with the leading exponent determined by the rate function.

The goal of our analysis is to develop a large deviation principle for our learning model, and use it to identify the direction and frequency of escapes. The difficulty in analyzing the evolution of beliefs comes from the interactions between the state $\xi_n$ in (8) and the beliefs $\gamma_n$ in (11). However we can decouple the analysis because the processes effectively operate on different scales. The state is a "fast" process with a variance independent of the gain $\varepsilon$, while the beliefs are a "slow" process whose variance decreases with $\varepsilon$. Thus in the next two sections we split the analysis into two steps reflecting the different scales.

### 6.2. Analyzing The State Evolution

As a first step, we analyze the fast state process, treating the slow belief process as fixed because its fluctuations are proportionately negligible. We later use the results from this stage to determine the large deviation properties of the slow belief process. In the convergence analysis in Section 5 we averaged over the variation in $\xi_n$, and showed the beliefs are typically driven by the mean dynamics $Eg(\gamma,\xi) = \bar{g}(\gamma)$. To analyze the contribution of fluctuations in $\xi_n$ to the escape dynamics, we now analyze events where $g$ differs from its expectation. In particular, we apply a large deviation principle due to Worms (1999) for the sample means of $g$, and we explicitly calculate the rate function.

For fixed beliefs $(R,\gamma)$, we now define a sequence of centered sample means for $g$. Let $g_i(\gamma,\xi)$ and $\bar{g}_i(\gamma)$, $i = 1,\ldots,n_y$ denote the columns of $g$ and $\bar{g}$, respectively. Then define:

$$X_{iT} = \frac{1}{T}\sum_{n=1}^{T} R^{-1} g_i(\gamma,\xi_n) - R^{-1}\bar{g}_i(\gamma), \tag{23}$$

and let $X_T = [X_{1T},\ldots,X_{n_yT}]$. Thus for a sample of size $T$, $X_T$ is the deviation of the belief updating term $R^{-1}g$ from its expectation. An important first step in our analysis is to characterize the large deviation properties of $X_T$.

A key component of our results is the Poisson equation (18) above, whose solution is $f(\gamma,\xi)$. As we discussed in Section 3.2, the solution of the Poisson equation is fundamental in studying additive functions of a Markov process. For example, Meyn and Tweedie (1993) provide a functional central limit theorem for processes like $X_T$ by using the Poisson equation to convert them into martingales. Thus the appearance of the Poisson equation here is rather natural. The following theorem, whose proof is in Appendix A.2.1, applies a result from Worms (1999) to characterize the large deviation properties of $X_T$.

Theorem 6.1. *Assume that Assumption 2.1 holds. Let $f(\gamma,\xi)$ be the solution of the Poisson equation (18), and define $Q(\gamma,R)$ as in (19) above. Then the sequence $\{X_T\}$ in (23) satisfies a large deviation principle with rate function:*

$$S_X(x) = \sup_{\alpha_1,\ldots,\alpha_{n_y}} \sum_{i=1}^{n_y} \left[\langle x,\alpha_i\rangle - \frac{1}{2}\alpha_i' Q(\gamma,R)\alpha_i\right]. \tag{24}$$

While it is not always easy to solve the Poisson equation, it admits an explicit solution in our model. In turn, this allows us to explicitly calculate the rate function $S_X$. To ease the presentation, we assume here that the state $\xi_n$ does not contain a constant term and the distribution $F$ of the shocks $W_n$ is the standard normal distribution $\Phi$. In this case, under Assumption 2.1 we see from (8) that the limit distribution $\pi$ is Gaussian with mean zero and variance $\Sigma(\gamma)$ which solves the Lyapunov equation:

$$\Sigma(\gamma) \;=\; \Theta(\gamma)\Sigma(\gamma)\Theta(\gamma)' + \Lambda\Lambda'. \tag{25}$$

By definition, $g$ in (9) is quadratic in $\xi$, so we can write each element of it as:

$$g_{ji}(\gamma,\xi) = \xi' V_{ji}(\gamma)\xi,\ i=1,\ldots,n_y,\ j=1,\ldots,n_x$$

for some matrices $V_{ji}$. We now show that the Poisson equation (18) reduces to a system of matrix Lyapunov equations, and that $Q$ is a fourth moment matrix of a Gaussian process. These simplifications allow us to apply our results in practice, as there are efficient numerical methods for Lyapunov equations. In Appendix A.2.3, we provide a proof and show how to alter the calculations when the state contains a constant.

Lemma 6.1. *Suppose that Assumption 2.1 holds and $F=\Phi$, the standard normal distribution. Then the solution of the Poisson equation (18) is a matrix-valued quadratic function with elements $f_{ji}(\gamma,\xi) = \xi' L_{ji}(\gamma)\xi$. The matrices $L_{ji}$ solve the Lyapunov equations:*

$$V_{ji}(\gamma) = L_{ji}(\gamma) - \Theta(\gamma)' L_{ji}(\gamma)\Theta(\gamma). \tag{26}$$

*Then $Q$ from (19) can be written $Q(\gamma,R) = R^{-1}\bar{Q}(\gamma)R^{-1}$, where the elements of $\bar{Q}$ are for $j,k=1,\ldots,n_x$ (suppressing dependence on $\gamma$):*

$$\begin{aligned}\bar{Q}_{jk} = \sum_{i=1}^{n_y} \quad & [2\,\mathrm{tr}(V_{ji}\Sigma V_{ki}\Sigma) + \mathrm{tr}(V_{ji}\Sigma)\,\mathrm{tr}(V_{ki}\Sigma) - \mathrm{tr}(L_{ji}\Lambda\Lambda')\,\mathrm{tr}(\Theta' L_{ki}\Theta\Sigma) \\ & - \mathrm{tr}(L_{ki}\Lambda\Lambda')\,\mathrm{tr}(\Theta' L_{ji}\Theta\Sigma) - \mathrm{tr}(L_{ki}\Lambda\Lambda')\,\mathrm{tr}(L_{ji}\Lambda\Lambda')]\,.\end{aligned}$$

Although the calculating $Q$ may appear rather involved, it is straightforward to implement. Moreover, when the model is static, as in the example in Section 4, then the Poisson equation is trivial and matters simplify further.

### 6.3. Analyzing the Beliefs

We now turn to the second step of the analysis, analyzing the large deviation properties of beliefs. As we mentioned in Section 3.2 above, our Theorem 3.1 builds on results in the literature. In particular, Freidlin and Wentzell (1998) characterized the most probable escape path for continuous time diffusions as a solution of a variational problem. Their results have been extended to discrete time stochastic approximation models such as ours by Dupuis and Kushner (1989). We first present their results, which are very general. Then we show how the results in the previous section lead to the simplified expressions given in Theorem 3.1 above.

We now define the key objects of the theory. Conditional on an arbitrary $\xi_0$, define:

$$H_i(\gamma, \alpha_i; R) = \lim_{T \to \infty} \frac{1}{T} \log E_{\xi_0} \exp \left\langle \alpha_i, \sum_{n=1}^{T} R^{-1} g_i(\gamma, \xi_n) \right\rangle, \; i = 1, \ldots, n_y \tag{27}$$

When $\xi_n$ is i.i.d., $H_i$ reduces to the log moment generating function of $R^{-1} g_i$. But in the more general case when $\xi_n$ is persistent, $H_i$ averages over the time dependence in the state and thus is a type of long-run log moment generating function. We then define the Legendre transform of the $H_i$ functions as:

$$L(\gamma, \beta; R) = \sup_{\alpha_1, \ldots, \alpha_{n_y}} \sum_{i=1}^{n_y} \left[ \langle \alpha_i, \beta_i \rangle - H_i(\gamma, \alpha_i; R) \right]. \tag{28}$$

The *action functional* is then defined over absolutely continuous trajectories $\gamma(\cdot)$ by:

$$S(T, \gamma(\cdot)) = \int_0^T L(\gamma(s), \dot{\gamma}(s); R(s)) ds \tag{29}$$

where $R(\cdot)$ solves (15). (We let $S = +\infty$ for trajectories that are not absolutely continuous.) Finally then, we define the variational problem:

$$\bar{S} = \inf_{\gamma(\cdot), T} S(T, \gamma(\cdot)) \tag{30}$$

subject to (15) and the boundary conditions (21). The following theorem is the basis for our Theorem 3.1, and it compiles and applies results from Dupuis and Kushner (1989), Kushner and Yin (1997), and Dembo and Zeitouni (1998). A proof is given in Appendix A.2.2. A similar theorem was also stated in Cho, Williams, and Sargent (2002), who applied our Theorem 3.1.

THEOREM 6.2. *Under the conditions of Theorem 3.1, its conclusions hold with $\bar{S}$ defined by (30), where $\gamma(\cdot)$ attains the minimum.*

While Theorem 6.2 offers a characterization of the large deviation properties of beliefs and the most probable escape path, it is difficult to derive useful insights from the minimization problem (30) itself. Additionally, because of the complicated nature of $H$ and $S$, analysis of the escape dynamics appears to be a daunting task. The next result, whose proof is in Appendix A.2.1, is a key contribution of this paper. It uses the results of the previous section to provide a simple expression for the $H$ function in our model.

LEMMA 6.2. *Under the assumptions of Theorem 3.1, the $H_i$ function defined in (27) is given by:*

$$H_i(\gamma, \alpha_i; R) = \left\langle \alpha_i, R^{-1}\bar{g}_i(\gamma) \right\rangle + \frac{1}{2}\alpha_i' Q(\gamma, R)\alpha_i. \tag{31}$$

This implies that $H_i$ in (27) essentially consists of a mean part (the first term) and a deviation from the mean (the second term) due to the stochastic shocks and the exponentiation. The mean part is relatively straightforward, reflecting the mean dynamics $R^{-1}\bar{g}$. The expression for the deviation from the mean comes from Theorem 6.1, by noticing that connection between the terms in (27) and $X_T$ from (23). Large deviation results apply here because the product inside the expectation in (27) becomes dominated by the extreme terms. The result (31) resembles the standard formula for the mean of a log-normal random variable, where $Q$ plays the role of a covariance matrix. This is consistent with our discussion above, where we noted that $Q$ acts as a generalized covariance matrix for the agent's beliefs.

The Legendre transform of the $H_i$ functions (31) is then simply:

$$L(\gamma, \dot{\gamma}; R) \;=\; \frac{1}{2}\sum_{i=1}^{n_y} \left(\dot{\gamma}_i - R^{-1}\bar{g}_i(\gamma)\right)' Q(\gamma, R)^{\dagger} \left(\dot{\gamma}_i - R^{-1}\bar{g}_i(\gamma)\right),$$

where we note that $Q$ may be singular and thus $Q^{\dagger}$ is its pseudoinverse. If we then define:

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

we justify the evolution equation (16) and can rewrite $L$ as a quadratic form in $v$. Using these simplifications, the abstract variational problem (30) becomes the simpler control problem (20), and our Theorem 3.1 follows.

### 6.4. Characterizing the Escape Dynamics

In this section, we characterize the solution of the minimum cost control problem (20). The deterministic control problem (20) can be solved either by applying a maximum principle or dynamic programming. The maximum principle allows for explicit expressions, but dynamic programming sometimes leads to more direct solutions. In Appendix A.2.4, we provide a result from Fleming and Soner (1993) which justifies the use of a maximum principle here. Thus the Hamiltonian for (20) with states $(\gamma, R)$, co-states $(\alpha, \lambda)$, and control $v$ is:

$$\begin{aligned}\mathcal{H}(\gamma, R, \alpha, \lambda) &= \inf_v \sum_{i=1}^{n_y} \left\{ \frac{1}{2} v_i' Q(\gamma, R)^\dagger v_i + \alpha_i \cdot \left( R^{-1} \bar{g}_i(\gamma) + v_i \right) \right\} + \lambda \cdot \left( \bar{M}(\gamma) - R \right) \\ &= \sum_{i=1}^{n_y} \left[ \left\langle \alpha_i, R^{-1} \bar{g}_i(\gamma) \right\rangle - \frac{1}{2} \alpha_i' Q(\gamma, R) \alpha_i \right] + \lambda \cdot \left( \bar{M}(\gamma) - R \right).\end{aligned}$$

Therefore, by taking derivatives of the Hamiltonian, we see that the most probable escape path solves the differential equations:

$$\begin{aligned}\dot{\gamma} &= R^{-1} \bar{g}(\gamma) - Q(\gamma, R)\alpha \\ \dot{R} &= \bar{M}(\gamma) - R \\ \dot{\alpha} &= -\alpha \cdot R^{-1} \frac{\partial \bar{g}(\gamma)}{\partial \gamma} + \frac{1}{2} \alpha' \frac{\partial Q(\gamma, R)}{\partial \gamma} \alpha - \lambda \cdot \bar{M}_\gamma(\gamma) \\ \dot{\lambda} &= -\mathcal{H}_R + \lambda \cdot I,\end{aligned} \tag{32}$$

subject to the boundary conditions (21), where $\mathcal{H}_R$ is the derivative of the Hamiltonian with respect to the $R$ matrix.

Alternatively, we can characterize the solution of (20) by dynamic programming methods.[17] We define a value function $\hat{S}(x, y)$ analogous to $\bar{S}$ in (20), but with the terminal condition in (21) changed to $\gamma(T) = x, R(T) = y$. Then $\hat{S}$ satisfies the following Hamilton-Jacobi partial differential equation:

$$\mathcal{H}(x, y, -\hat{S}_x(x, y), -\hat{S}_y(x, y)) = 0. \tag{33}$$

Then $\bar{S} = \inf_{x,y} \hat{S}(x, y)$ subject to $x \in \partial G$.

In some special cases, the PDE (33) is explicitly solvable. But in general we must rely on numerical solutions, and the ODEs (32) are typically easier to implement numerically. Our results in Section 4 relied on the numerical solution of this system of differential equations. Given an initial condition for the co-states $(\alpha(0), \lambda(0))$, it is easy to integrate the ODEs until $\gamma$ hits the boundary of the set $G$. This determines $T, \gamma(\cdot)$ and allows us to evaluate the action functional $S$ as in (29). We then solve the minimization problem (30) by minimizing over $(\alpha(0), \lambda(0))$.

## 7. CONCLUSION

In this paper we have analyzed two sources of dynamics that govern adaptive learning models: mean dynamics which pull an agent's beliefs toward a limit point, and escape dynamics which push them away. We have provided a precise characterization of these dynamics, and we have illustrated how they can arise in an example. We have shown that as the gain decreases to zero (across sequences), the beliefs converge

[17] See Dembo and Zeitouni (1998), Exercise 5.7.36.

in a weak sense to a self-confirming equilibrium. However ongoing stochastic shocks may occasionally lead beliefs to escape from the self-confirming equilibrium. We developed new theoretical methods to characterize the escape dynamics, and showed how to apply them in a simple economic model. There we saw how a misspecificaton may generate locally self-reinforcing dynamics, which lead to recurrent large deviations from the self-confirming equilibrium.

The methods that we have developed have potentially broad applications. The results in this paper have already been applied to analyze recurrent inflations and stabilizations (by Cho, Williams, and Sargent (2002) among others), deflationary liquidity traps (by Bullard and Cho (2005)), currency crises (by Cho and Kasa (2008)), and fluctuations in prices resembling price wars (by Williams (2009)). Further, our analysis is not limited to learning models. The same methods can be applied to filtering and recursive estimation problems, which could have interesting implications for the performance of estimators. However, absent in standard estimation problems is the feedback between estimates and observations which drove much of our results. Learning models thus provide a natural framework in which escape dynamics can play an important role.

# APPENDIX A

## A.1. CONVERGENCE RESULTS

For these results, we stack $(\gamma, R)$ into a vector $\theta$ and write (11)-(12) compactly as:

$$\theta_{n+1} = \theta_n + \varepsilon b(\theta_n, \xi_{n+1}).$$

Then we define $\bar{b}(\theta) = Eb(\theta, \xi_n)$ and $v_{n+1} = b(\theta_n, \xi_{n+1}) - \bar{b}(\theta_n)$. The following are the necessary assumptions for Theorem 5.1 above. We include an $\varepsilon$ superscript on $\theta_n^\varepsilon$ and $v_n^\varepsilon$ to emphasize their dependence on the gain.

Assumptions A.1.

1. *The random sequence* $\{\theta_n^\varepsilon; \varepsilon, n\}$ *is tight.*[1]
2. *For each compact set* $A$, $\left\{b(\theta_n^\varepsilon, \xi_{n+1})1_{\{\theta_n^\varepsilon \in A\}}; \varepsilon, n\right\}$ *is uniformly integrable.*[2]
3. *The ODE* $\dot{\theta} = \bar{b}(\theta)$ *has a point* $\bar{\theta}$ *which is asymptotically stable.*[3]
4. *The function* $\bar{b}(\theta)$ *is continuous.*

[1] A random sequence $\{X_n\}$ is *tight* if

$$\lim_{K \to \infty} \sup_n P\left(|X_n| \geq K\right) = 0.$$

[2] A random sequence $\{X_n\}$ is *uniformly integrable* if

$$\lim_{K \to \infty} \sup_n E\left(|X_n|\, 1_{\{|X_n| \geq K\}}\right) = 0.$$

[3] A point $\bar{x}$ 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 $\varepsilon > 0$ such that if $|x(0) - \bar{x}| \leq \varepsilon$, then $|x(t) - \bar{x}| \leq \delta$ for all $t$.

*5. For each $\delta > 0$, there is a compact set $A_\delta$ such that* $\inf_{n,\varepsilon} P(v_n^\varepsilon \in A_\delta) \geq 1 - \delta$.

*Proof* (Theorem 5.1). The result follows directly from Theorem 8.5.1 in Kushner and Yin (1997). The theorem requires their additional assumptions (A8.1.9), (A8.5.2), (A8.5.3) and (A8.5.5) which hold trivially here, since $Eb(\theta_n^\varepsilon, \xi_{n+1}) = \bar{b}(\theta_n^\varepsilon)$ is independent of $\xi_{n+1}$. This implies that the limit in (A8.1.9) is identically zero and that the $\beta_n^\varepsilon$ terms in (A8.5.2) and (A8.5.5) are also identically zero. Further, their conditions (A8.5.1) and (A8.5.3) are then equivalent and given by part 2 above. The theorem is also stated under a weaker condition which is implied by part 3 above. ∎

Note that parts 1, 2, and 5 of Assumptions A.1 hold in our model with i.i.d. Gaussian shocks $W_n$. (The conditions are even easier to verify with bounded shocks.) The tightness in part 1 follows because for each $\theta$, $b(\theta, \xi_{n+1})$ is a quadratic function of standard normal random variables. Therefore $P(|b(\theta, \xi_{n+1})| \geq K) = f(K)$ for some function $f$ that goes to zero as $K$ goes to infinity. Since the one-step transitions satisfy this property, any finite number of steps does as well. Further, since the property holds for all $\theta$, we have that $P(|\theta_n^\varepsilon| \geq K) \to 0$ as $K \to \infty$, and so the sequence is tight. For part 2, note that $|b(\theta_n, \xi_{n+1})|^2$ consists of normally distributed random variables up to the fourth order, and so has finite expectation, which implies the uniform integrability. Finally part 5 holds because $v_n$ consists of normally distributed random variables up to the second order, and thus can be bounded to arbitrary accuracy on an appropriate compact set. The remaining conditions 3 and 4 must be verified in particular settings.

## A.2. LARGE DEVIATIONS

### A.2.1. Analyzing the State Evolution

Here we collect the proofs of our results in Section 6.2.

*Proof* (Theorem 6.1). We first state the assumptions underlying the result of Worms (1999) that we use in Theorem 6.1, then we show that in our model Assumption 2.1 suffices for them to hold. The assumptions are stated in terms of a functional autoregressive model that generalizes (8):

$$\xi_{n+1} = f(\xi_n) + \sigma(\xi_n) W_{n+1},$$

where $\xi \in \mathbb{R}^{n_s}$, $f : \mathbb{R}^{n_s} \to \mathbb{R}^{n_s}$, $\sigma : \mathbb{R}^{n_s} \to \mathbb{R}^{n_s \times n_w}$, and $n_s = n_y + n_w$.

ASSUMPTIONS A.2.

*1. For some norm $|\cdot|$ on $\mathbb{R}^{n_s}$ and $\mathbb{R}^{n_s \times n_w}$, $f$ and $\sigma$ satisfy the following relation for $\beta > 2$ and for any $x, y \in \mathbb{R}^{n_s}$ :*

$$|f(x) - f(y)| + E(|W|^\beta)^{\frac{1}{\beta}} |\sigma(x) - \sigma(y)| \leq \alpha |x - y|,$$

*where* $0 < \alpha < 1$.

*2. For some $a \geq 0$, $b > 0$ with $1 \leq a + b \leq \beta/2$, with $\|\cdot\|$ denoting the Euclidean norm, and for any $\gamma \in \mathcal{G}$ the function $g(\gamma, \cdot)$ satisfies:*

$$\|g(\gamma, x) - g(\gamma, y)\| \leq C \|x - y\|^a \left( \|x\|^a + \|y\|^b + 1 \right).$$

Note that in the linear-quadratic model given by (8), $f(\xi) = \Theta\xi$ and $\sigma(\xi) = \Lambda$. Therefore under our Assumption 2.1, part 1 of Assumptions A.2 holds for any $\beta$ such that $E(|W|^\beta) < \infty$, with the norm $|\cdot|$ being a standard matrix norm. Further, since $g(\gamma, \xi)$ is quadratic in $\xi$, part 2 of Assumptions A.2 holds with (for example) $a = b = 1$. Then Theorem 6.1 follows from Theorem 2 of Worms (1999). ∎

*Proof* (Lemma 6.2). This result relies on an important duality between exponential limits and large deviations due to Varadhan. The following is Theorem 4.3.1 from Dembo and Zeitouni (1998), adapted to the general setting of Section 6.1.

Theorem A.1. *Assume that the sequence $\{X_T\}$ satisfies a large deviation principle on $\mathcal{X}$ with the convex rate function $S_X$. Also suppose that for some $\lambda > 1$ we have:*

$$\limsup_{T\to\infty} \frac{1}{T} \log E \exp\left(\lambda T \langle \alpha, X_T \rangle\right) < \infty.$$

*Then we have that:*

$$\lim_{T\to\infty} \frac{1}{T} \log E \exp\left(T \langle \alpha, X_T \rangle\right) = \sup_{x\in\mathcal{X}} \left\{ \langle \alpha, x \rangle - S_X(x) \right\}.$$

We now use this result to prove Lemma 6.2. In Theorem 6.1 we establish the rate function for the $\{X_T\}$ process (23), which is a partial sum of a functional autoregressive process. The rate function (24) depends on the matrix-valued function $Q(\gamma, R)$, and can be written as:

$$S_X(x) = \frac{1}{2} x' Q(\gamma, R)^\dagger x.$$

Theorem A.1 applies, as the rate function is clearly convex, and the moment condition holds (since $g$ is quadratic function of normally distributed random variables). Thus via the convex duality of $S_X$ and $Q$ we get (31). ∎

### A.2.2. Analyzing the Beliefs

This section collects assumptions and proofs for the results in Section 6.3. First we state the additional assumptions necessary for the large deviation principle.

Assumptions A.3.

*1a. The sequence $\{|R_n^{-1} g(\gamma_n, \xi_{n+1})|\}$ is almost surely bounded by some constant $K < \infty$.*

*1b. There exist a $\sigma-$algebra $\mathcal{F}_n \supset \sigma(\gamma_i, i \leq n)$ and constants $\kappa > 1, B < \infty$ such that for all $n$ and $s \geq 0$:*

$$P\left( \left| R_n^{-1} g(\gamma_n, \xi_{n+1}) \right| \geq s |\mathcal{F}_n \right) \leq B \exp(-s^\kappa) \ a.s.$$

For Assumptions A.3, we require either part 1a or 1b to hold.

*Proof* (Theorem 3.1). The results follow directly from Theorem 6.2, with the simplifications made by Theorem 6.1 and Lemma 6.2 in Section 6.2. ▮

*Proof* (Theorem 6.2). (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 a stability condition satisfied by part 3 of Assumptions A.1. Their assumption 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 1b of Assumptions A.3, 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.

(2): The result is an application of Kushner and Yin (1997) Theorem 6.10.1, whose assumptions follow directly under the boundedness condition of 1a of Assumptions A.3. The identification of the $H$ function follows from Dupuis and Kushner (1989), Theorems 4.1 and 5.3.

(3): Kushner and Yin (1997) establish an upper bound on mean escape times in Theorem 6.10.6. After establishing part 2 of Theorem 6.2, the results here follow directly from Dembo and Zeitouni (1998), Theorem 5.7.11.

(4): The first part also follows from Theorem 5.7.11 in Dembo and Zeitouni (1998), which is analogous to Theorem 2.1 in Freidlin and Wentzell (1998). The second part follows from Theorem 2.3 in Freidlin and Wentzell (1998). Our phrasing of the result follows Dupuis and Kushner (1987). ▮

### A.2.3. Solving the Poisson Equation

This section proves our result on the solution of the Poisson equation in Section 6.2, then discusses the necessary modifications when the state $\xi$ includes a constant term.

*Proof* (Lemma 6.1). For simplicity, we suppress dependence on $\gamma$. Since $g$ is quadratic and $\xi$ is Gaussian, we guess that $f$ is also quadratic with elements $f_{ji}$:

$$f_{ji}(\xi) = \xi' L_{ji} \xi, \tag{A.1}$$

for some matrices $L_{ji}$. Taking conditional expectations gives:

$$f_{ji}(\xi) - E[f_{ji}(\xi_{n+1})|\xi_n = \xi] = \xi' L_{ji} \xi - \xi' \Theta' L_{ji} \Theta \xi - \mathrm{tr}(L_{ji} \Lambda\Lambda').$$

Additionally, we know that the elements of $g$ satisfy:

$$g_{ji}(\xi) - \bar{g}_{ji} = \xi' V_{ji} \xi - \mathrm{tr}(V_{ji} \Sigma).$$

Thus for $f$ in (A.1) to solve (18), the $L_{ji}$ matrices must satisfy the Lyapunov equations given in (26). This verifies our guess of the form of $f$ and also insures that:

$$\mathrm{tr}(L_{ji} \Lambda\Lambda') = \mathrm{tr}(V_{ji} \Sigma),$$

so that the constant terms in the Poisson equation match.

With the Poisson solution $f$, we can then calculate the matrix-valued $Q$ function as in (19). We factor this function as: $Q(\gamma, R) = R^{-1}\bar{Q}(\gamma)R^{-1}$. Then using the solution for $f$ above, we have that the $(j,k)$ element of the $\bar{Q}$ matrix is of the form:

$$\bar{Q}_{jk} = \sum_{i=1}^{n_y} E\left[(\xi' L_{ji}\xi)(\xi' L_{ki}\xi) - (\xi'\Theta' L_{ji}\Theta\xi + \operatorname{tr}(L_{ij}\Lambda\Lambda'))(\xi'\Theta' L_{ki}\Theta\xi + \operatorname{tr}(L_{ki}\Lambda\Lambda'))\right] \tag{A.2}$$

where the expectation is with respect to the stationary distribution $\pi$. Note that (A.2) is the expectation of a quartic function of normal random variables. Evaluating the expectation gives the expression in the text. ∎

We now extend Lemma 6.1 to the case where the state $\xi_n$ contains a constant. It is useful here to isolate the constant, so we define $\xi_n = [\tilde{\xi}_n', 1]'$ and use (8) to write:

$$\tilde{\xi}_{n+1} = \Theta_0 + \underline{\Theta}\tilde{\xi}_n + \underline{\Lambda} W_{n+1} \tag{A.3}$$

where the matrices $(\Theta_0, \underline{\Theta}, \underline{\Lambda})$ are drawn from the matrices $(\Theta, \Lambda)$. Thus $\tilde{\xi}_n$ is Gaussian with mean $\widetilde{\mu}$ and covariance $\widetilde{\Sigma}$ given by:

$$\begin{aligned} \widetilde{\mu} &= (I - \underline{\Theta})^{-1}\Theta_0 \\ \widetilde{\Sigma} &= \underline{\Theta}\widetilde{\Sigma}\underline{\Theta}' + \underline{\Lambda}\,\underline{\Lambda}'. \end{aligned}$$

The distribution of the full vector $\xi_n$ is thus partially degenerate. The invariant measure $\pi$ is Gaussian, with mean $\mu$ and covariance $\Sigma$ given by:

$$\mu = \begin{bmatrix} \widetilde{\mu} \\ 1 \end{bmatrix}, \; \Sigma = \begin{bmatrix} \widetilde{\Sigma} & 0 \\ 0 & 0 \end{bmatrix}.$$

The function $g$ is still quadratic in $\xi$, and thus we can write its elements as:

$$g_{ji}(\xi) = \xi' \widetilde{V}_{ji} \xi$$

for some matrices $\widetilde{V}_{ji}$, as above. However for some calculations it is now useful to renormalize $g$ by making it a function of a mean zero process. So we define a process that centers the variable components of $\xi$ but retains the constant term: $z_n = [\tilde{\xi}_n' - \widetilde{\mu}', 1]'$. This process is also Gaussian with the linear evolution:

$$z_{n+1} = \begin{bmatrix} \underline{\Theta} & 0 \\ 0 & 1 \end{bmatrix} z_n + \Lambda W_{n+1} = \widetilde{\Theta} z_n + \Lambda W_{n+1},$$

where $\underline{\Theta}$ is as in (A.3) and $\Lambda$ is from (8). Then we rewrite $g$ as a function of the new $z$ process:

$$g_{ji}(z) = z' V_{ji} z + k_{ji},$$

where we partition $\widetilde{V}_{ji}$ conformably to $z$ and define:

$$
\begin{aligned}
V_{ji} &= \begin{bmatrix} \widetilde{V}_{ji,11} & \widetilde{V}_{ji,12} + \widetilde{V}_{ji,11}\mu \\ \widetilde{V}_{ji,21} + \mu'\widetilde{V}_{ji,11} & 0 \end{bmatrix}, \\
k_{ji} &= \widetilde{V}_{ji,22} + \mu'\widetilde{V}_{ji,11}\mu + \mu'\widetilde{V}_{ji,12} + \widetilde{V}_{ji,21}\mu.
\end{aligned}
$$

With these modifications, the same calculations as in Lemma 6.1 imply that the solution $f$ to the Poisson equation (18) takes the form:

$$f_{ji}(z) = z'L_{ji}z + \hat{f}_{ji},$$

where the matrices $L_{ji}$ satisfy the Lyapunov equations parallel to (26):

$$V_{ji} = L_{ji} - \widetilde{\Theta}'L_{ji}\widetilde{\Theta}.$$

The constants $\hat{f}_{ji}$ are undetermined, but this is inessential as they drop out of the matrix $Q$.

We can then evaluate $Q$ as in (A.2) above with $z$ replacing $\xi$, and thus we gain a few additional terms due to the means. For example, the first product term in (A.2) is now (the second term is similar):

$$
\begin{aligned}
E\left[(z'L_{ji}z)(z'L_{ki}z)\right] = 2\,\mathrm{tr}(L_{ji}\Sigma L_{ki}\Sigma) + \mathrm{tr}(L_{ji}\Sigma)\,\mathrm{tr}(L_{ki}\Sigma) + e_s'L_{ji}'\Sigma L_{ki}e_s \\
+ e_s'L_{ki}\Sigma L_{ji}e_s + e_s'L_{ji}\Sigma L_{ki}e_s + e_s'L_{ji}\Sigma L_{ki}'e_s,
\end{aligned}
$$

where $e_s$ is a $n_s$ vector of all zeros except for a one in the final position (to select the constant).

### A.2.4. Characterizing the Escape Dynamics

The following lemma from Fleming and Soner (1993) justifies the use of a maximum principle in solving the control problem (20).

LEMMA A.1. *Under our standing assumptions, $L(\gamma, \beta)$ in (28) is strictly convex in $\beta$ and obeys a superlinear growth condition:*

$$\frac{L(\gamma, \beta)}{|\beta|} \to \infty \text{ as } |\beta| \to \infty.$$

*Therefore the solution of the problem (20) can be determined by the differential equations (32).*

*Proof.* Fleming and Soner (1993) establish that if $H$ is strictly convex in $\alpha$, then $L$ is strictly convex in $\beta$, and also that if $H$ is superlinear in $\alpha$, then $L$ is superlinear in $\beta$. We have shown in (31) that $H$ is a quadratic function of $\alpha$ and so is strictly convex and superlinear. The conclusion follows from results in Section I.6 and I.8 of Fleming and Soner (1993). ∎

## A.3. VERIFYING THE ASSUMPTIONS IN THE EXAMPLE

In this section, we formally verify the necessary conditions of our theorems above for our example. Since the model is static and there is no time dependence in the beliefs, Assumption 2.1 is immediately satisfied as long as there is a finite solution to the firm's problem. Thus we require that the firm's perceived demand curve slope downward, and to simplify our results below we assume that the slope is bounded away (negatively) from zero. Similarly for prices to be positive, we require that the firm's intercept be positive. Thus we take the feasible set to be:

$$\mathcal{G} = \{\gamma : \gamma_0 \geq 0,\ \gamma_1 \leq \delta < 0\}.$$

Notice that the self-confirming equilibrium we identify in the text is strictly within this set (as long as $b_1 < \delta < 0$ and $b_0 > 0$), and the escape sets $G$ we analyze are within this set as well.

Then we need to verify Assumptions A.1 and A.3. Following our discussion after the proof of Theorem 5.1 above, we know that parts 1, 2, and 5 of Assumptions A.1 hold. Further, part 1b of Assumptions A.3 is immediate since we assume that the shocks are Gaussian. Since we consider an algorithm with $R_n$ fixed, part 4 of Assumptions A.1 simply requires the continuity of $\bar{g}(\gamma)$. From the expressions in Section 4.2 we see clearly that $\bar{g}(\gamma)$ is continuous on $\mathcal{G}$, so part 4 of Assumptions A.1 holds.

The only remaining condition is part 3 of Assumptions A.1, the asymptotic stability of the ODE. Note again that for the algorithm here we can consider just the ODE for $\gamma$. We have identified the self-confirming equilibrium $\bar{\gamma}$ above, which is the unique equilibrium point of the ODE. Further, one can show that the eigenvalues of the Jacobian matrix of $\bar{g}$ evaluated at $\bar{\gamma}$ all have strictly negative real parts, so that it is locally asymptotically stable. Global stability is more difficult to establish explicitly. However numerical analysis of the ODE suggests that (at least for the parameterizations we consider) the ODE is in fact asymptotically stable on $\mathcal{G}$.

## REFERENCES


1. Bullard, James and In-Koo Cho (2005) "Escapist Policy Rules," *Journal of Economic Dynamics and Control*, 29: 1841-1865.
2. Cho, In-Koo and Kenneth Kasa (2008) "Learning Dynamics and Endogenous Currency Crises," *Macroeconomic Dynamics*, 12: 257-85.
3. Cho, In-Koo, Noah Williams, and Thomas J. Sargent (2002) "Escaping Nash Inflation," *Review of Economic Studies*, 69: 1-40.
4. Dembo, Amir and Ofer Zeitouni (1998) *Large Deviations Techniques and Applications,* Second Edition, Springer-Verlag, New York.
5. Dupuis, Paul and Harold J. Kushner (1987) "Stochastic Systems with Small Noise, Analysis and Simulation," *SIAM Journal on Applied Mathematics,* 47: 643-661.
6. Dupuis, Paul and Harold J. Kushner (1989) "Stochastic Approximation and Large Deviations: Upper Bounds and w.p.1. Convergence," *SIAM Journal on Control and Optimization,* 27: 1108-1135.
7. Ellison, Martin and Tony Yates (2007) "Escaping Volatile Inflation," *Journal of Money, Credit, and Banking,* 39: 981-993.
8. Evans, George and Seppo Honkapohja (2001) *Learning and Expectations in Macroeconomics,* Princeton University Press.

9. Evans, George, Seppo Honkapohja, and Noah Williams (2009) "Generalized Stochastic Gradient Learning," *International Economic Review,* forthcoming.
10. Fleming, Wendell and H. Mete Soner (1993) *Controlled Markov Processes and Viscosity Solutions,* Springer-Verlag, New York.
11. Foster, Dean and H. Peyton Young (1990) "Stochastic Evolutionary Game Dynamics," *Theoretical Population Biology,* 38: 219-232.
12. Freidlin, Mark I. and Alexander D. Wentzell (1998) *Random Perturbations of Dynamical Systems,* 2nd edition. Springer-Verlag, New York.
13. Fudenberg, Drew and David Levine (1998) *The Theory of Learning in Games,* MIT Press, Cambridge.
14. Kandori, Michihiro, George Mailath and Rafael Rob (1993) "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica,* 61: 29-56.
15. Kasa, Kenneth (2004) "Learning, Large Deviations, and Recurrent Currency Crises," *International Economic Review*, 45: 141-173.
16. Kreps, David (1998) "Anticipated Utility and Dynamic Choice," in *Frontiers of Research in Economic Theory* (D. P. Jacobs, E. Kalai, and M. I. Kamien, eds.), Cambridge University Press.
17. Kushner, Harold J. and George G. Yin (1997) *Stochastic Approximation Algorithms and Applications,* Springer-Verlag, New York.
18. Maier, Robert S. and Daniel L. Stein (1997) "Limiting Exit Distributions in the Stochastic Exit Problem," *SIAM Journal on Applied Mathematics,* 57: 752-790.
19. Marcet, Albert and Thomas J. Sargent (1989) "Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models," *Journal of Economic Theory,* 48: 337-368.
20. Marcet, Albert and Juan Pablo Nicolini (2003) "Recurrent Hyperinflations and Learning," *American Economic Review*, 93: 1476-1498.
21. McGough, Bruce (2006) "Shocking Escapes," *Economic Journal*, 116: 507-528.
22. Meyn, Sean P. and Richard L. Tweedie (1993) *Markov Chains and Stochastic Stability,* Springer-Verlag, New York.
23. Sargent, Thomas J. (1999) *The Conquest of American Inflation,* Princeton University Press.
24. Sargent, Thomas J. and Noah Williams (2003) "Impacts of Priors on Convergence and Escapes from Nash Inflation," *Review of Economic Dynamics*, 8: 360-391.
25. Sargent, Thomas J., Noah Williams, and Tao Zha (2006). "Shocks and Government Beliefs: The Rise and Fall of American Inflation," *American Economic Review*, 96: 1193-1224.
26. Sargent, Thomas J., Noah Williams, and Tao Zha (2008). "The Conquest of South American Inflation," working paper, NYU and University of Wisconsin - Madison.
27. Williams, Noah (2002) "Stability and Long Run Equilibrium in Stochastic Fictitious Play," working paper, University of Wisconsin - Madison.
28. Williams, Noah (2009) "Equilibria, Escapes, and Price Wars," working paper, University of Wisconsin - Madison.
29. Worms, Julien (1999) "Moderate Deviations for Stable Markov Chains and Regression Models," *Electronic Journal of Probability,* 4(8): 1-28.
30. Young, H. Peyton (1993) "The Evolution of Conventions," *Econometrica*, 61: 57-84.