

Escape Dynamics in Learning Models

Noah Williams*

Department of Economics, University of Wisconsin - Madison

E-mail: nmwilliams@wisc.edu

Revised October 5, 2016

This paper illustrates and characterizes how adaptive learning can lead to recurrent large fluctuations. 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, generating large movements in observed outcomes. I characterize the escape dynamics by drawing on the theory of large deviations, developing new results which make this theory directly applicable in a class of learning models. The likelihood, frequency, and most likely direction of escapes are all characterized by a deterministic control problem. I illustrate my results with two simple examples.

1. INTRODUCTION

In this paper I show how learning dynamics can provide an important mechanism generating recurrent large fluctuations in economic models. There is by now a substantial literature on adaptive learning in economics, with important work both in macroeconomics (see Evans and Honkapohja, 2001 and 2009) and game theory (see Fudenberg and Levine, 1998 and Young, 2004). 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 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 I develop and apply methods to characterize these rare departures. Following Sargent (1999), I call these large deviations *escape dynamics*.

In particular, I focus on situations in which agents are uncertain of their economic environment. These agents make forecasts and base their actions on subjective models, which I assume are linear regressions that they update as they observe data. Agents allow for structural change in their environment, which leads them to discount past data. I also allow agents' subjective models to be misspecified, and thus following Sargent (1999) my 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 Philipp Kircher and three anonymous referees for useful comments. I also thank Fernando Alvarez, Jim Bullard, Marco Cagetti, Jeff Campbell, Xiaohong Chen, In-Koo Cho, Amir Dembo, John Duffy, Dana Heller, Ken Kasa, Andrew Postlewaite, Juha Seppälä, Ted Temzelides, Harald Uhlig, and especially Lars Peter Hansen and Thomas Sargent for helpful comments, discussions, and suggestions.

¹See 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 I 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. I show that agents' beliefs converge to a differential equation, called 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). My 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, which in turn causes them to change their behavior. This affects the data they observe, and hence can feed back on their beliefs. In this process, agents 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 small gains. But for any positive gain, escapes may occasionally recur, and when they do, they have a very particular form. To analyze the escape dynamics, I draw on the theory of large deviations. When agents' beliefs escape, with high probability they closely follow a deterministic path called the *most probable escape path*. I show how to find this path, and characterize the likelihood and frequency of escapes.

While my 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 infrequent aberrations, and my results characterize these rare tail events. However, Sargent (1999) demonstrated that escape dynamics may be an important force generating fluctuations in some classes of models. In this paper I provide two examples illustrating these different types of applications.

First, I study a learning variant of a Lucas (1978) asset pricing model. Previous work by Timmermann (1993), Branch and Evans (2011), Benhabib and Dave (2014), and Adam, Marcet, and Nicolini (2016) has shown how learning can generate additional asset price volatility, booms and busts, and predictability in returns, all of which can improve empirical performance. I follow Benhabib and Dave (2014), who use a model where agents forecast future asset prices based on their perceived relation between prices and dividends, and show that beliefs under learning have a fat-tailed distribution. I apply my results to determine the rate at which large belief fluctuations decline with decreasing gain, and show that the fluctuations are asymmetric. For small gain, it is increasingly likely that large fluctuations in prices are triggered by perceived increases (rather than decreases) in the responsiveness of prices to fundamentals.

My second example shows how escape dynamics may be the most striking feature of some environments. Escapes can lead to episodes which look like switches in equilibria or changes in regime, which would not occur absent learning. As in the work of Sargent (1999), Cho, Williams, and Sargent (2002), McGough (2006), Ellison and Yates (2007), Ellison and Scott (2013), and Williams (2016), the model I study has a unique self-confirming equilibrium, but the time series show repeated, regular escapes. As in the asset pricing model, the escapes are less frequent with smaller gains, but unlike that

model they retain the same regular character and size. In particular, I study a model of a monopolist learning its demand curve, subject to both cost and demand shocks. 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 are not present when the cost and demand shocks are independent. My results explain this difference, and characterize both the frequency and direction of escapes. I describe how the escapes are due to locally self-reinforcing dynamics which arise once stochastic shocks push beliefs away from the SCE, a feature which is shared by the other models in the literature.

Overall, my main contribution is the derivation of a simple deterministic control problem whose solution characterizes the escape dynamics in a class of belief-dependent linear dynamic models. My results build on Dupuis and Kushner (1989), who 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. My results simplify the general theory and make it directly applicable for learning in linear models.

There are two sources of dynamics which govern agents' beliefs. The mean dynamics govern the expected behavior of beliefs and drive the convergence toward the SCE. Escape dynamics are driven by unlikely shock realizations, and I show that they can be interpreted as a perturbation of the mean dynamics. My 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. I then apply standard control theory methods to characterize the solution of the cost minimization problem.

While the mean dynamics and convergence have been well-studied in the literature, until recently 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 analyzed games with multiple equilibria, using large deviation methods to determine the stochastic transition rates between equilibria. Although my results can be used to analyze multiplicities, as in Williams (2014), here I focus on models with a unique equilibrium.

As mentioned above, Sargent (1999) introduced escape dynamics and large deviation theory for learning models, and provided much of the motivation for this paper. Since drafts of this paper were first circulated, the results developed here have been applied and extended in a variety of settings. My results were first applied by Cho, Williams, and Sargent (2002) to analyze Sargent's (1999) model. Further related papers which build on or apply my results include Sargent and Williams (2005), Bullard

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

³Important papers in this literature include Kandori, Mailath, and Rob (1993) and Young (1993). See Section 5.3 below for more discussion.

and Cho (2005), Ellison, Graham, and Vilmunen (2006), McGough (2006), Ellison and Yates (2007), Cho and Kasa (2008), Williams (2009), Ellison and Scott (2013), and Kolyuzhnov, Bogomolova, and Slobodyan (2014). See Section 5.3 for more discussion. This later literature either applies the results in this paper, or studies special cases or approximations of my results. None provide as thorough or general a characterization of the escape dynamics.

The rest of the paper is organized as follows. In Section 2, I introduce the baseline model, and discuss the equilibrium concept and learning formulation. In Section 3, I establish the convergence of beliefs, while Section 4 provides the large deviation results characterizing the escape dynamics. Section 5 then describes and analyzes the examples discussed above, and relates the results in this paper to the rest of the literature. The appendix collects proofs and statements of technical results.

2. THE MODEL

In this section I describe the class of models studied in the paper. I focus on linear models in which agents forecast or make decisions based on estimated models which they update over time. For simplicity, I focus on a single agent setting.⁴

2.1. The Basic Setup

Time is discrete $n = 0, 1, 2, \dots$, and there is state vector $y_n \in \mathbb{R}^{n_y}$, which follows a belief-dependent linear stochastic process:

$$y_{n+1} = \bar{A}(\gamma)y_n + \bar{\Sigma}(\gamma)W_{n+1} \quad (1)$$

where $\gamma \in \mathbb{R}^{n_x}$ is a vector of beliefs or regression coefficients, and $W_{n+1} \in \mathbb{R}^{n_w}$ is an i.i.d. shock vector with distribution F . I mostly focus on the case where F is Gaussian, but some results are stronger for bounded shocks. As I show below, the specification (1) is relatively flexible, representing a reduced form where the true state is dependent on agent's beliefs, either through expectations or decisions.

The agent is uncertain about his environment, with γ parameterizing his subjective beliefs which he updates over time as he observes data. I adopt a relatively flexible specification for the agent's beliefs which permits some parts of the environment to be known with certainty, but yet allows the subjective model to be misspecified. I focus on the case where the agent learns about a single equation in the state evolution, but the results generalize to multiple equations without difficulty except in notation. I parse the state vector $y_n = [z_n, c_n']'$ where z_n is the scalar state whose evolution is uncertain, while the agent knows the evolution of c_n . For example, if the state y_n includes lagged dynamics, then c_n would be the lags, whose dynamics are known. I also allow the agent's model to be misspecified, in that it may omit some relevant variables. Thus instead of conditioning on the full state vector y_n he may only consider a sub-vector $x_n \in \mathbb{R}^{n_x}$.

⁴Williams (2016) studies dynamic games using the same approach. But strategic interaction raises some additional issues which are not essential here.

The subjective model is then:

$$z_{n+1} = \gamma' x_n + \eta_{n+1}. \quad (2)$$

Here η_{n+1} is a vector of regression errors which are believed to be orthogonal to the regressors x_n :

$$\tilde{E}[x_n(z_{n+1} - \gamma' x_n)'] = 0. \quad (3)$$

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.

2.2. Structures Leading to the Reduced Form

Two different but related structural formulations lead to the belief-dependent linear model in the reduced form (1): linear expectational models and linear state space models. These classes of models are widely used in applications, and below I study an example of each.

Linear expectational models often arise in linearizing the equilibrium conditions of nonlinear dynamic models. They can be written in the form:

$$A_0 y_n = A_1 y_{n-1} + B \tilde{E}_n y_{n+1} + \Sigma W_n. \quad (4)$$

By appropriately defining the matrices (A_0, A_1, B, Σ) this form can include both exogenous dynamics (without expectational terms) and endogenous dynamics. Linear rational expectations models, which take this form and identify \tilde{E} with the mathematical expectations operator, have a long history in economics and are widely used in applied work. But here \tilde{E} defines the agent's subjective expectations as in equations (2)-(3). As above, I assume that expectations only matter for the sub-vector z_n and so simplify:

$$B \tilde{E}_n y_{n+1} = B_0 \tilde{E}_n z_{n+1} = B_0 \gamma x_n = B_0 \gamma K_x y_n$$

where the first equality is by assumption, the second uses the subjective beliefs, and the third defines $x_n = K_x y_n$. In the terminology of Evans and Honkapohja (2001), equation (2) is the perceived law of motion (PLM), while substituting the expectations into (4) gives the actual law of motion (ALM) for y_n :

$$[A_0 - B_0 \gamma K_x] y_n = A_1 y_{n-1} + \Sigma W_n.$$

Assuming that the leading matrix is invertible and shifting forward one period, we get a version of the belief-dependent reduced form (1) with $\bar{A}(\gamma) = [A_0 - B_0 \gamma K_x y_n]^{-1} A_1$ and $\bar{\Sigma}(\gamma) = [A_0 - B_0 \gamma K_x y_n]^{-1} \Sigma$.

Linear state space models are a related structure which place more emphasis directly on decisions. Suppose an agent chooses actions a_n to affect the state y_n according to:

$$y_{n+1} = A y_n + B a_n + \bar{\Sigma} W_{n+1}. \quad (5)$$

Such models often arise as approximations by linearizing decision problems. The agent does not know the full state evolution (5), but instead acts based on his subjective model (2) which we now write as:

$$z_{n+1} = \gamma'_y y_n + \gamma'_a a_n + \eta_{n+1},$$

which emphasizes that the agent may be uncertain about the impact of his actions. Thus $\gamma = [\gamma'_y, \gamma'_a]'$ and $x_n = [y'_n, a'_n]'$, and to allow for misspecification some elements of γ_y may be zero. If the agent has quadratic intertemporal preferences over y_n and a_n , then by standard results his optimal choice is a linear decision rule:

$$a_n = h(\gamma)y_n, \tag{6}$$

where I emphasize the dependence on the beliefs γ . Substituting the rule (6) in the state evolution (5) gives a version of the reduced form (1) with $\bar{A}(\gamma) = A + Bh(\gamma)$. To implement this structure over time, I follow much of the learning literature in adopting the Kreps (1998) and Sargent (1999) “anticipated utility” approach: at every date the agent chooses his actions to maximize his utility given his current beliefs, then he updates beliefs based on observed outcomes.

2.3. Self-Confirming Equilibrium

Following Fudenberg and Levine (1998) and Sargent (1999), I now define a self-confirming equilibrium as a matrix of beliefs which are consistent with the agent’s observations. First, let $\xi_{n+1} = [y'_n, W'_{n+1}]'$ and define g as the function whose expectation is zero in (3):

$$g(\gamma, \xi_{n+1}) = x_n(z_{n+1} - \gamma'x_n)' \tag{7}$$

Note that x_n and y_{n+1} are both linear functions of ξ_{n+1} , and hence so is z_{n+1} . Thus g is a quadratic function of ξ_{n+1} , a particular structure I will exploit below.

The key orthogonality condition (3) can then be written as $\bar{E}g(\gamma, \xi_{n+1}) = 0$. In a self-confirming equilibrium this orthogonality condition holds under the objective probability measure induced by (1) as well. That is, the agent’s beliefs are confirmed by his observations. For the objective expectation to make sense, I assume that given γ , y_n has a stationary distribution denoted π . I later constrain the evolution of beliefs to ensure that π exists. Thus define \bar{g} as the unconditional expectation of g :

$$\bar{g}(\gamma) = E[g(\gamma, \xi_{n+1})] = \int \int g(\gamma, y, W) d\pi(y) dF(W). \tag{8}$$

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

2.4. Adaptation

So far the agent's beliefs have been fixed, and now I turn to how he updates them with observations. I specify that the agent learns via a constant gain recursive least squares algorithm:

$$\gamma_{n+1} = \gamma_n + \varepsilon R_n^{-1} g(\gamma_n, \xi_{n+1}) \quad (9)$$

$$R_{n+1} = R_n + \varepsilon \phi(x_n x_n' - R_n). \quad (10)$$

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. I introduce the factor ϕ in (10) to also include the class of generalized stochastic gradient learning rules, as studied by Evans, Honkapohja, and Williams (2010). 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 (9)-(10) would be a recursive representation of the standard OLS estimator. 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 nonstationary environments, and are 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, I assume that y_n in (1) is stationary. If the agent were to know the true model (5), then standard assumptions would ensure this. However since the agent's model (2) differs from the truth and his beliefs evolve over time, we need to constrain the learning rule so that it does not induce instability in the state evolution. The following assumption restricts beliefs γ to guarantee the existence of a unique, stable invariant distribution π for y_n .

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

One way to ensure this stability in practice is to impose a projection facility on (9), as in Marcet and Sargent (1989) which restricts the updating rule so that the estimates stay in the set. I do not explicitly deal with such a facility here, as I assume that the SCE $\bar{\gamma}$ is in interior of \mathcal{G} and analyze escapes to points that remain in the interior of \mathcal{G} . While I focus on dynamic models, my results also hold for static models, where instability is not an issue. In that case I allow y_n to include a constant, which implies that \bar{A} has one unit eigenvalue.⁶

⁵Sargent and Williams (2005) discuss the performance of the constant gain algorithm for drifting coefficients. Evans, Honkapohja, and Williams (2010) show that constant gain rules are robust to misspecification.

⁶Even in the dynamic case, when y_n contains a constant term, only $n_y - 1$ eigenvalues of the state matrix in (1) need be less than one.

Finally for some purposes it will be convenient to stack the beliefs into a single vector $\theta_n = [\gamma'_n, \text{col}(R_n)']'$. Then we can write (9)-(10) as:

$$\theta_{n+1} = \theta_n + \varepsilon \Psi(\theta_n, \xi_{n+1}) \quad (11)$$

where Ψ is a quadratic function of ξ_{n+1} , and we define $\bar{\Psi}(\theta) = E\Psi(\theta, \xi_{n+1})$.

3. CONVERGENCE OF BELIEFS

I first show that on average, the agent will be drawn toward a self-confirming equilibrium, then later I characterize escapes from the SCE. The results in this section follow from Kushner and Yin (1997), and are analogous to results in Evans and Honkapohja (2001). Marcet and Sargent (1989) first used differential equations to characterize the limits of learning models, but they focused on decreasing gain models and so had a different characterization of convergence.

3.1. Overview

All of the results in the paper consider small gain limits. In time series the gain is constant, so this limit looks across different series indexed by the gain. I emphasize this by writing γ_n^ε . As $\varepsilon \rightarrow 0$ 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}^ε :

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

Then we can re-write (9) as:

$$\frac{\gamma_{n+1}^\varepsilon - \gamma_n^\varepsilon}{\varepsilon} = (R_n^\varepsilon)^{-1} \bar{g}(\gamma_n^\varepsilon) + v_{n+1}^\varepsilon. \quad (12)$$

Note that (12) is similar to a finite-difference approximation of a time derivative, on a time scale where ε is the increment between observations. Letting $\varepsilon \rightarrow 0$, this approximation becomes arbitrarily good. Along this same limit, a law of large numbers ensures that v_n^ε converges to zero. Thus in the limit we obtain the differential equations:

$$\dot{\gamma} = R^{-1} \bar{g}(\gamma) \quad (13)$$

$$\dot{R} = \phi[\bar{M}(\gamma) - R] \quad (14)$$

Equation (14) carries out a similar limit for (10), where we use the notation $\bar{M}(\gamma) = E(x_n x_n')$. I call these ODEs the *mean dynamics*, as they govern the expected evolution of the agent's beliefs. Theorem 3.1 below makes this formal. Note that an equilibrium point $\bar{\gamma}$ of (13) is a self-confirming equilibrium, and let $\bar{M}(\bar{\gamma}) = \bar{R}$. Thus if the SCE is stable under the ODE, as $\varepsilon \rightarrow 0$ the agent's beliefs (9)-(10) converge to $(\bar{\gamma}, \bar{R})$.

3.2. Formal Results

To proceed with the analysis more formally, I define a time scale to convert the discrete time belief evolution into a continuous time process. Let the ε be the continuous time interval and interpolate between the discrete iterations in the learning rule (9)-(10):

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

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 $\varepsilon \rightarrow 0$ 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 $\varepsilon \rightarrow 0$ the interpolated processes converge to the solution of the ODEs derived informally in (13)-(14).

In Appendix A.1 I 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 I show there that many of the conditions are satisfied in the baseline model. The remaining conditions require that $\bar{g}(\gamma)$ and $\bar{M}(\gamma)$ be continuous and that the system of ODEs (13)-(14) have an asymptotically stable point $(\bar{\gamma}, \bar{R})$. I show below how to verify these conditions in practice.

THEOREM 3.1. *Under Assumptions A.1, as $\varepsilon \rightarrow 0$, $(\gamma^\varepsilon(\cdot), R^\varepsilon(\cdot))$ 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 $\varepsilon \rightarrow 0$ over time ($t \rightarrow \infty$) agents' beliefs converge weakly to the SCE. The same limiting ODE characterizes decreasing gain algorithms, such as the recursive least squares algorithm studied by Marcet and Sargent (1989) which sets $\varepsilon = \frac{1}{n+1}$ in (9)-(10). But with decreasing gain the beliefs typically converge with probability one as $n \rightarrow \infty$, while we obtain weak convergence with constant gain. The weaker notion of convergence here means that for any given gain ε , occasional departures from the SCE may persist over time.

4. 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 ε . However, for a fixed $\varepsilon > 0$ we may observe such rare escapes, and large deviation theory allows us to characterize them. I now show that when escapes happen, they are very likely to happen in a particular most probable way with a calculable frequency.

4.1. Large Deviations and Escape Dynamics

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.

First I define the key objects of interest. To simplify the presentation, I initialize all paths at the SCE.⁷ The results use the same time scale as (12) where ε is the time increment between n and $n + 1$, and the stacked beliefs θ_n as in (11). Define $\bar{\theta} = [\bar{\gamma}', \text{col}(\bar{R})']'$, and recalling that \mathcal{G} is the set of stable beliefs γ , let $\theta = [\gamma', \text{col}(R)']' \in \mathcal{T}$ if $\gamma \in \mathcal{G}$.

DEFINITION 4.1. Fix an $\varepsilon > 0$, a time horizon $\bar{n} < \infty$ (which may depend on ε), and a compact set $G \subset \mathcal{T}$ with non-empty interior and $\bar{\theta} \in G$. Let $\theta^\varepsilon(t), t \in [0, \bar{n}\varepsilon]$ be the piecewise linear interpolation of $\{\theta_n^\varepsilon\}$.

1. An *escape path* from G is a sequence $\{\theta_n^\varepsilon\}_{n=0}^{\bar{n}}$ solving (11) such that $\theta_0^\varepsilon = \bar{\theta}$ and $\theta_m^\varepsilon \notin G$ for some $m \leq \bar{n}$. Let $\Gamma^\varepsilon(G, \bar{n})$ be the set of escape paths.
2. For any sequence $\{\theta_n^\varepsilon\}_{n=0}^{\bar{n}}$ solving (11) with $\gamma_0 = \bar{\gamma}$, define the (first) *escape time* from G as:

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

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

For small gains, any path $\{\theta_n^\varepsilon\}$ spends most of its time near the SCE $\bar{\theta}$, and if noise pushes it away, it tends to be drawn back. While with unbounded shocks 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 μ_2 neighborhood of the SCE does not return to a smaller μ_1 neighborhood of it. That is, once a regular escape path starts to escape, it does not turn back.

My key results characterize bounds on the probability of escape, the mean escape time, and the *most probable escape path*.⁸ I show 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 .

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

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

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

$$\dot{\gamma} = R^{-1}\bar{g}(\gamma) + v_\gamma \quad (15)$$

$$\dot{R} = \phi[\bar{M}(\gamma) - R] + v_R. \quad (16)$$

Note the similarity of (15) to (12): for the mean dynamics the perturbation vanishes, but it resurfaces to govern the escape dynamics. That is, the perturbation v_n^ε in (12) has a zero mean, so it typically becomes negligible for small gains and the beliefs track the mean dynamics. Unlikely sequences of shocks lead to an escape, which I analyze by introducing the perturbations (v_γ, v_R) in (15)-(16). Alternative escape paths are associated with alternative perturbations, and I evaluate the likelihood of alternative escape paths by a cost function which penalizes more unlikely perturbations.

I now work with the composite beliefs θ , collecting the perturbations from (15)-(16) into $v = [v'_\gamma, \text{col}(v_R)']'$. 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 the beliefs were normally distributed, then their variance would be a natural measure of volatility. However recall that g is a quadratic function of a linear stochastic process, so we need a measure of its volatility which appropriately captures the tail behavior of beliefs. For a general i.i.d. random variable, we can summarize its distribution by its moment generating function, or equivalently the log of it, which is also known as the cumulant generating function. For an i.i.d random variable ξ_{n+1} , define the log moment generating function of beliefs for $\alpha \in \mathbb{R}^{n_\theta}$ as:

$$\begin{aligned} H(\theta, \alpha) &= \log E \exp \langle \alpha, \Psi(\theta, \xi_{n+1}) \rangle \\ &= \langle \alpha, \bar{\Psi}(\theta) \rangle + \log E \exp \langle \alpha, v \rangle, \end{aligned} \quad (17)$$

where the second line uses (15)-(16), and $\langle \cdot, \cdot \rangle$ denotes an inner product. Define the Legendre transform of this function as:

$$\begin{aligned} L(\theta, v) &= \sup_{\alpha} [\langle \alpha, \bar{\Psi}(\theta) + v \rangle - H(\theta, \alpha)] \\ &= \sup_{\alpha} [\langle \alpha, v \rangle - \log E \exp \langle \alpha, v \rangle]. \end{aligned} \quad (18)$$

The function L plays the role of the instantaneous cost function for the belief perturbations v . As Dembo and Zeitouni (1998) emphasize, it captures the tail behavior of the distribution of $\Psi(\theta, \xi_{n+1})$, which is crucial for analyzing the escape dynamics. If v were distributed normally with mean zero and variance Σ_v we would have:

$$H(\theta, \alpha) = \langle \alpha, \bar{\Psi}(\theta) \rangle + \frac{1}{2} \alpha' \Sigma_v \alpha, \quad \text{and:} \quad L(\theta, v) = \frac{1}{2} v' \Sigma_v^{-1} v.$$

The cost function L would weight the perturbations v by the covariance matrix, just as we discussed above. However in our case v is a quadratic form, so the calculations are a bit more complex, even when ξ_n is i.i.d.

When, as in our baseline model, ξ_n is not i.i.d., we need a more general cost function for perturbations. This is based on taking H to be the long run moment generating function of ξ_n . Conditional on an arbitrary ξ_0 , define:

$$H(\theta, \alpha) = \lim_{T \rightarrow \infty} \frac{1}{T} \log E_{\xi_0} \exp \left\langle \alpha, \sum_{n=1}^T \Psi(\theta, \xi_n) \right\rangle. \quad (19)$$

This more general H averages over the temporal dependence in ξ_n . Once again we define L as its Legendre transform as in the first line of (18), which is the cost function in the dynamic model. While calculating it is not always an easy task, Section 4.3 below shows how to explicitly compute H and thus L in our model.

4.2. Characterizing the Escape Dynamics

We analyze escapes on a fixed continuous time horizon $\bar{T} < \infty$, and set $\bar{n} = \bar{T}/\varepsilon$. Thus $\bar{n} \rightarrow \infty$ as $\varepsilon \rightarrow 0$. To characterize escapes from a set G , we choose the perturbations v in (15) which push beliefs to the boundary ∂G in the most cost effective way:

$$\bar{S} = \inf_{v(\cdot), T} \int_0^T L(\theta(s), v(s)) ds \quad (20)$$

where the minimization is subject to (15), (16) and:

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

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 $\bar{\theta}$ to the boundary of G .

The following theorem shows how the control problem (20) characterizes escapes. It compiles and applies results from Dupuis and Kushner (1989), Kushner and Yin (1997), and Dembo and Zeitouni (1998). Below I derive the explicit form of the control problem in (20) for our class of models, which makes the theory directly applicable. We fix a set G and horizon \bar{T} as above, and recall that τ^ε is an escape time, with the escape taking place at $\theta^\varepsilon(\tau^\varepsilon)$. 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.2 are in Appendix A.1 and A.2, respectively. A proof is given in Appendix A.2. A similar result was stated in Cho, Williams, and Sargent (2002), who applied our Theorem 4.1.¹⁰

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

¹⁰Cho, Williams, and Sargent (2002) used the results from an earlier version of this paper, which had a slightly different characterization of the rate function \bar{S} . The previous results improperly used a theorem of Worms (1999) to simplify the calculation of H , but this result does not apply in our setting. In practice, this only matters for dynamic models, not the static model Cho, Williams, and Sargent (2002) focused on where ξ_n is i.i.d. In that paper the results for the static model used direct calculations as we do below.

THEOREM 4.1. *Suppose that Assumptions 2.1, A.1, and A.2 hold, let $\theta^\varepsilon(\cdot)$ be the piecewise linear interpolation of $\{\theta_n^\varepsilon\}$, and let $\theta(\cdot) : [0, \bar{T}] \rightarrow \mathbb{R}^{n_\theta}$ solve (20).*

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

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

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

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

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

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

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

$$\lim_{\varepsilon \rightarrow 0} P(\|\theta^\varepsilon(\tau^\varepsilon) - \theta(T)\| < \delta | \{\theta_n^\varepsilon\} \in \Gamma^\varepsilon(G, \bar{n})) = 1.$$

$$\text{Moreover: } \lim_{\varepsilon \rightarrow 0} P(\|\theta^\varepsilon(t) - \theta(t)\| < \delta, t < \tau^\varepsilon(\{\theta_n^\varepsilon\}) | \{\theta_n^\varepsilon\} \in \bar{\Gamma}^\varepsilon(G, \bar{n})) = 1.$$

Proof. See Appendix A.2. **■**

Part (1) shows that the probability of observing an escape on a bounded time interval is exponentially decreasing in the gain ε , with the rate given by the minimized cost function \bar{S} . The next three parts establish stronger results for bounded shocks.¹¹ Part (2) shows that in this case the asymptotic inequality in part (1) becomes an equality. Part (3) shows that for small ε the escape times from the SCE become close to $\exp(\bar{S}/\varepsilon)$.¹² 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.

4.3. Calculating the Cost Function

While Theorem 4.1 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 (20) itself. Additionally, because of the complicated nature

¹¹Although 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.

¹²But notice that δ is fixed, and thus as $\varepsilon \rightarrow 0$ the interval around $\exp(\bar{S}/\varepsilon)$ expands.

of H and \bar{S} , analysis of the escape dynamics appears to be a daunting task. In this section I provide a simple expression for the H function in our model with normally distributed shocks. This allows me to implement the theory in practice.

First, recall that $\Psi(\theta, \xi_{n+1})$ is a quadratic function of $\xi_{n+1} = [y'_n, W'_{n+1}]'$. Then we can write the following:

$$\langle \alpha, \Psi(\theta, \xi_n) \rangle = -\frac{1}{2} (y'_n V_{yy} y_n + 2y'_n V_{yw} W_{n+1} + W'_{n+1} V_{ww} W_{n+1})$$

for some matrices V_{yy}, V_{yw}, V_{ww} , where V_{yy} and V_{ww} are symmetric. Clearly the matrices depend on the beliefs θ as well as the vector α .

I first focus on the simpler case where $\xi_n = [1, W'_n]$, which happens for instance in static models (and is used in an example below). In this case we only need to compute the simpler version of H in (17). The proof of this result is in Appendix A.3.

LEMMA 4.1. *Assume that $W_n \sim N(0, I)$ and $\xi_n = [1, W'_n]$. Then under the assumptions of Theorem 4.1, the H function defined in (17) is given by:*

$$H(\theta, \alpha) = -\frac{1}{2} \log |V_{ww} + I| + \frac{1}{2} (V_{yw} (V_{ww} + I)^{-1} V'_{yw} - V_{yy}) \quad (22)$$

The H function in the static case is given by a relatively simple expression depending on the V matrices, but this can imply some rather complex dependence on the underlying beliefs θ . In the fully dynamic model, we need to compute the more complex H function in (19) above, which averages over the temporal dependence in ξ_n . The next result, whose proof is in Appendix A.3, is a key contribution of this paper. For this result we need some standard conditions from linear optimal control theory, controllability and detectability, which us allow to calculate the limit in (19).¹³

LEMMA 4.2. *Assume that $W_n \sim N(0, I)$ and that the assumptions of Theorem 4.1 hold. In addition, suppose that $(\bar{A}, \bar{\Sigma}, V_{yw})$ is controllable and detectable for all $\theta \in G$. Then the H function defined in (19) is given by:*

$$H(\theta, \alpha) = -\frac{1}{2} \log |V_{ww} + \bar{\Sigma}' \Theta \bar{\Sigma} + I|, \quad (23)$$

where Θ solves the discrete algebraic Riccati equation:

$$\Theta = \bar{A}' \Theta \bar{A} + V_{yy} - (V_{yw} + \bar{A}' \Theta \bar{\Sigma}) (V_{ww} + \bar{\Sigma}' \Theta \bar{\Sigma} + I)^{-1} (V_{yw} + \bar{A}' \Theta \bar{\Sigma})'. \quad (24)$$

The general form of the H function in (23) is the same as in the static model (22), with the first term capturing the variance of the beliefs. In the dynamic case, this is

¹³Kwakernaak and Sivan (1972) say that the n -dimensional system (A, B, C) is controllable if the column vectors of $(B, AB, A^2B, \dots, A^{n-1}B)$ span the whole n -dimensional space. The definition of detectability is more involved (see p. 465), but a sufficient condition is that the row vectors of $[C', (CA)', (CA^2)', \dots, (CA^{n-1})']'$ span the whole n -dimensional space.

a long-run variance, computed by solving a Riccati equation of the same form that appears in linear-quadratic control problems. The second term in the static expression (22) represents the contribution of the constant. The corresponding component in the dynamic case would capture the effects of the initial condition y_0 , but such effects vanish once we take the limit in (19).

4.4. Solving the Minimization Problem

The deterministic control problem (20) can be solved in a standard way by applying a maximum principle. The minimized Hamiltonian for (20) with state θ , co-state λ , and control $\beta = \bar{\Psi}(\theta) + v$ is:

$$\begin{aligned}\mathcal{H}(\theta, \lambda) &= \inf_{\beta} \{L(\theta, \beta) + \langle \lambda, \beta \rangle\} \\ &= \inf_{\beta} \left\{ \sup_{\alpha} [\langle \alpha, \beta \rangle - H(\theta, \alpha)] + \langle \lambda, \beta \rangle \right\} \\ &= -H(\theta, -\lambda),\end{aligned}$$

where the second equality carries out the minimization and maximization, noting the convex duality of L and H . If we parameterize $\alpha = -\lambda$, then by taking derivatives of the Hamiltonian, we see that the most probable escape path solves the differential equations:

$$\begin{aligned}\dot{\theta} &= H_{\alpha}(\theta, \alpha) \\ \dot{\alpha} &= -H_{\theta}(\theta, \alpha)\end{aligned}\tag{25}$$

subject to the boundary conditions (21).

In some special cases, the differential equations are explicitly solvable, but in general we must rely on numerical solutions. Given an initial condition for the co-states $\alpha(0)$, it is easy to integrate the ODEs until θ hits the boundary of the set G or reaches the terminal time \bar{T} . Paths which do not escape are assigned arbitrarily large cost values. For paths which do escape, this procedure determines $T, \theta(\cdot)$ and we can evaluate the cost function $\int_0^T L(\theta(s), v(s))ds$ as in (20). We then solve the minimization problem (20) by minimizing over the initial co-state $\alpha(0)$.

5. EXAMPLES

I now discuss two simple examples which illustrate escape dynamics in the two different frameworks discussed above: expectational and state space models. The first is a simple asset pricing model, where learning gives rise to occasional large fluctuations in beliefs and asset prices. I show that for smaller gains these fluctuations become less frequent, and they also become asymmetric. My results describe both the frequency and the asymmetry of the escapes. The second example shows how escape dynamics may be a dominant feature of a mode, as in other applications of escape dynamics discussed in Section 5.3 below. The model has a unique self-confirming equilibrium

which is stable under learning. In one parameterization, the model behaves like the first example, with escapes being rare large movements away from the SCE. However under a different parameterization, the model has recurrent escapes from the SCE which are a prominent feature of the time series and always lead in the same direction. My results characterize the escapes and explain these differences.

5.1. An Asset Pricing Model

5.1.1. Setup

My first example is an asset pricing model that was analyzed by Benhabib and Dave (2014). The model illustrates the role of expectations in a simple univariate environment, focusing on agents' forecasts of future prices. Benhabib and Dave (2014) used different but related large deviation results to show that the tail of the belief distribution under learning follows a power law, and thus is subject to large fluctuations.¹⁴ Relative to their paper, I analytically characterize the rate of fluctuations and asymmetries in the distribution of beliefs, and show that the predictions explain behavior observed in simulations.

The model can be derived by log-linearizing the equilibrium conditions of a Lucas (1978) asset pricing model. The exogenous driving process in the model is the scalar dividend x_n , which follows a stationary AR(1) process:

$$x_{n+1} = \rho x_n + \sigma W_{n+1}.$$

The asset price, which we denote z_n , is determined by the linear expectational equation:

$$z_n = b\tilde{E}_n z_{n+1} + ax_n, \quad (26)$$

with the current price determined by dividends and expectations of future prices. However the agent does not know (26) but instead has the subjective model as in (2):

$$z_{n+1} = \gamma x_n + \eta_{n+1}.$$

This is a simple special case of the linear expectational model in Section 2.2 above, with $y_n = [z_n, x_n]$. Using the subjective model for expectations gives the actual law of motion for the price:

$$z_{n+1} = (b\gamma + a)\rho x_n + (b\gamma + a)\sigma W_{n+1}.$$

Together with the law of motion for dividends, this yields the belief-dependent evolution (1). However since the asset price z_n is a linear function of x_n and W_{n+1} and has no dynamics of its own, we do not need to track it, so can simplify and define $\xi_{n+1} = [x_n, W_{n+1}]'$.

As in Benhabib and Dave (2014), I assume that agents learn with a constant gain stochastic gradient rule. Thus I set $\phi = 0$ in (10) and use the constant weight $R = 1$.

¹⁴Benhabib and Dave (2014) apply results for the limiting distribution of scalar multiplicative processes. Thus they consider large time limits which hold for any gain setting, but only apply for this class of scalar models which are linear in beliefs. By contrast, my results are small gain limits and apply more broadly to multidimensional models that are nonlinear in beliefs, but are only approximate for small gains.




FIGURE 1. Simulated price series for different gain settings. The red dashed lines show the SCE mean price of zero.

The key belief function g is then:

$$g(\gamma, \xi_{n+1}) = ((b\rho - 1)\gamma + a\rho) x_n^2 + (b\gamma + a)\sigma x_n W_{n+1}.$$

Taking (unconditional) expectations gives the mean dynamics \bar{g} :

$$\bar{g}(\gamma) = ((b\rho - 1)\gamma + a\rho) \frac{\sigma^2}{1 - \rho^2}.$$

This function is linear, hence the unique SCE is $\bar{\gamma} = a\rho/(1 - b\rho)$ is globally stable, and it is also the rational expectations equilibrium. Since there is no misspecification in this model, the agent could eventually learn to forecast rationally. As the dynamics only come through the exogenous dividends x_n , Assumption 2.1 ensuring stability is satisfied for $0 \leq \rho < 1$.

Figure 1 plots three representative simulated time series of prices z_n from the model for different values of the gain ε .¹⁵ In each case, the price fluctuates around the SCE value of zero, but there are periods of large booms and crashes. For smaller gain settings, the price volatility declines (note the scale of the vertical axis) and the extreme events become less frequent. Figure 2 plots the simulated beliefs γ_n^ε underlying these price series. The large spikes in asset prices (either positive or negative) correspond to large movements in beliefs. For smaller gain settings, beliefs spend more time near the SCE value (shown with red dashed lines), just as the results suggest. In addition, large asset price movements tend to be driven by movements away from the SCE in the positive direction, reflecting an over-estimation of the relationship between prices and dividends.

¹⁵I choose parameters in line with Benhabib and Dave's (2014) estimates: $\rho = 0.9$, $\sigma = 0.2$, $b = 0.9$ and $a = (1 - b - 3)\rho + 3 = 0.39$.




FIGURE 2. Simulated slope coefficient series γ_n . The red dashed lines show the SCE slope.




FIGURE 3. Left panel: Escape distribution for the slope coefficient ($\gamma^\epsilon(\tau^\epsilon)$) for different gains (ϵ). Right panel: 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).

5.1.2. Escape Dynamics

I now apply my results to characterize the large movements in beliefs and prices as due to escape dynamics. The key terms in the calculations are:

$$V_{yy} = -2\alpha((b\rho - 1)\gamma + a\rho), \quad V_{yw} = -\alpha(b\gamma + a)\sigma, \quad V_{ww} = 0$$

In addition, since we only need track the scalar dynamics of x_n , the Riccati equation (24) is simply a quadratic equation in Θ whose solution is:

$$\Theta = \frac{\rho^2 - 1 + \sigma^2 V_{yy} - 2\rho\sigma V_{yw} + \sqrt{(\rho^2 - 1 + \sigma^2 V_{yy} - 2\rho\sigma V_{yw})^2 - 4\sigma^2(V_{yw}^2 - V_{yy})}}{2\sigma^2}.$$




FIGURE 4. Simulated and predicted regular escape paths with $\varepsilon = \frac{1}{300}$. Shown are the path closest to the mean time (blue solid line) and 90% band (black dashed lines) of the simulated distribution along with the predicted escape path (red dashed lines).

I then solve the control problem (20), obtaining the minimized cost \bar{S} and the most probable escape path $\gamma(\cdot)$. For the escape set G , I choose the interval $|\gamma_n^\varepsilon - \bar{\gamma}| < 1.5$, and similar results held for intervals of different length.

Figure 3 illustrates the escape distribution as well as the predictions from Theorem 4.1. In solving (20), I find that the most probable path exits the interval in the positive direction, $\gamma(T) = \bar{\gamma} + 1.5$. Part (4) of the theorem says that as $\varepsilon \rightarrow 0$ the distribution of escape points should concentrate at this point. This is documented in the left panel of Figure 3, which shows a histogram of the slope coefficient at the time of escape ($\gamma^\varepsilon(\tau^\varepsilon)$) from 20,000 simulated escapes for three different gain settings. For larger gains, the escape distribution is relatively symmetric, with escapes at both ends of the interval. However as the gain decreases, the escapes become increasingly concentrated at the upper end, as the theorem predicts. In addition, part (3) of Theorem 4.1 predicts that for small gain the mean escape times (on the continuous time scale $\tau^\varepsilon = n\varepsilon$) increase exponentially in $1/\varepsilon$ at rate \bar{S} . The right panel of Figure 3 plots the log mean escape times from the simulations along with our prediction, and bands covering 90% of the simulated distribution. Note that theorem only predicts the slope of the line shown, giving the exponential rate of increase.¹⁶ For relatively large gains, the escapes occur more rapidly than the results suggest, but for small gains the predictions match the simulations quite well.

Finally, part (4) of Theorem 4.1 states that in the limit all regular escapes remain close to the most probable path. I define regular paths by setting $\mu_1 = 0.01$. That is, I run 20,000 simulations which terminate when the beliefs are 1.5 units from the

¹⁶The 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.

SCE $\bar{\gamma}$. Many of these escape paths have small sojourns away from $\bar{\gamma}$ before returning to a neighborhood of it. I extract the regular part of the path by finding the last time the beliefs were within 0.01 units of the SCE, and keep the escape path from that time forward. Figure 4 summarizes the escape path distribution, showing the most probable path resulting from our calculation (20), a path from the simulations with length closest to the mean, and paths with lengths corresponding to the 5% 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 period close to the SCE, followed by a rapid increase in the slope. The shape of all the regular paths is quite similar, and the predicted most probable escape path is quite close to the mean from the simulations. Thus my results accurately predict the entire time path of beliefs during an escape.

5.2. A Monopoly Model

The previous example illustrated how escape dynamics can generate large fluctuations around a self confirming equilibrium. I now show how escape dynamics in learning models can generate time series which differs qualitatively from the SCE, with recurrent, regular periods of extended deviations from the SCE. I study a model of a monopoly firm facing an unknown linear demand curve, which is subject to a shock. The firm produces at a constant expected cost per unit in each period, with the realized cost depending on a shock. The correlation between cost and demand shocks is a key parameter. For simplicity, the model is static, with the only dynamics coming through learning.

5.2.1. Setup

The state vector consists of output z_n and a constant: $y_n = [z_n, 1]'$, and the shock vector $W_n = [W_{1n}, W_{2n}]'$ is a 2 dimensional standard normal random vector. For simplicity I set the expected marginal cost to zero, so effectively the firm chooses its markup a_n over marginal cost, with the cost shock determining the realized price p_n :

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

Output is given by the static linear demand curve:

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

Note the dating convention: z_{n+1} is the current period's output, which depends on the current price p_n . When $\rho \neq 0$, the shocks to costs and demand are correlated. This is a special case of the state space model in Section 2.2 above.

The firm does not know its demand curve, but instead sets its price based on its subjective model (2), which here takes the form:

$$z_{n+1} = \gamma_0 + \gamma_1 p_n + \eta_{n+1} \tag{27}$$

thus $x_n = [1, p_n]'$. The firm continues to believe that the error η_{n+1} is orthogonal to the regressors x_n , but when $\rho \neq 0$, the belief equation (27) is misspecified as p_n and

z_{n+1} are driven by correlated shocks. The firm maximizes expected profits based on (27), which can be written as:

$$\tilde{E}_n[p_n z_{n+1}] = \tilde{E}_n[\gamma_0 p_n + \gamma_1 p_n^2],$$

where I use (3). Since there are no dynamics in the model, static profit maximization problem gives the optimal markup as in (6):

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

As above, I assume the firm learns with a generalized stochastic gradient rule. The key function $g = [g_1, g_2]'$ from (7) is given explicitly here by:

$$\begin{aligned} g_1(\gamma, \xi_n) &= b_0 - \gamma_0 - (b_1 - \gamma_1) \frac{\gamma_0}{2\gamma_1} + (b_1 + \rho - \gamma_1) \sigma_a W_{2n} + \sigma_y W_{1n} \\ g_2(\gamma, \xi) &= g_1(\gamma, \xi) \left(-\frac{\gamma_0}{2\gamma_1} \right) + \left(b_0 - \gamma_0 - (b_1 - \gamma_1) \frac{\gamma_0}{2\gamma_1} \right) \sigma_a W_{2n} + \\ &\quad (b_1 + \rho - \gamma_1) \sigma_a^2 W_{2n}^2 + \sigma_a \sigma_y W_{1n} W_{2n}. \end{aligned}$$

Since $\xi_n = [1, W_{1n}, W_{2n}]'$ is i.i.d., taking expectations gives $\bar{g} = [\bar{g}_1, \bar{g}_2]'$ as in (8):

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

Since the model is static, Assumption 2.1 which ensures stability is satisfied.¹⁷ 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) . In the SCE the monopolist over-estimates the slope of his demand curve, and this perceived higher demand elasticity makes him set a lower price than would be optimal if he knew the demand curve.

The results above show that as $\varepsilon \rightarrow 0$ the beliefs converge to the SCE $\bar{\gamma}$.¹⁸ Thus we expect that for small ε , beliefs will remain near the SCE $\bar{\gamma}$, and the price will exhibit small fluctuations around the expected price $h(\bar{\gamma})$. However we find that the model exhibits striking behavior, as shown in Figures 5 and 6, which plot some simulated outcomes from the model for different settings of ρ . Figure 5 plots simulated time

¹⁷Some 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.

¹⁸Appendix A.4 provides formal detail and verifies all necessary assumptions for the example.




FIGURE 5. Simulated price series p_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 6. Simulated slope coefficient series γ_{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 $\rho = -1$ ($\bar{\gamma}_1 = b_1 + \rho = -2$).




FIGURE 7. 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).

paths of prices p_n for $\rho = 0$ and $\rho = -1$, while Figure 6 plots the estimated slope coefficients γ_{1n} from the same simulations.¹⁹ Each figure also plots two different gain settings when $\rho = -1$. When $\rho = 0$, and thus the model (27) is correctly specified, the price and belief series behave as expected. The price and slope coefficients are near the SCE levels throughout, with movements away following no regular pattern.

But when $\rho = -1$, there are recurrent episodes in which the firm raises its price sharply, during 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. I now show that this is precisely what my results predict.

5.2.2. Escape Dynamics

In solving the control problem (20) to characterize the escapes, I 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 increases in the slope, while $G(-r)$ is the half-ball with decreases in slope. The minimized cost \bar{S} as a function of r is shown in Figure 7 for $\rho = 0$ and $\rho = -1$. When $\rho = 0$, the cost is symmetric, increasing rapidly in both directions. Thus escapes become quite rare for small gains, and are equally likely to result in increases or decreases of the slope

¹⁹The other parameters in the model are: $b_0 = 10$, $b_1 = -1$, $\sigma_a = \sigma_y = 0.2$.




FIGURE 8. Escape distributions for the slope coefficient ($\gamma_1^\epsilon(\tau^\epsilon)$) for different gains (ϵ), different radii (r), and two settings of ρ . Each panel considers a different (ρ, r) combination and plots the distribution for $\epsilon = 0.1$ (blue bars) and $\epsilon = 0.01$ (red bars). The black dotted lines show the terminal point of the most probable escape ($\gamma_1(T)$).

coefficient. But when $\rho = -1$ the minimized cost function increases rapidly in the negative direction, but in the positive direction it has a long, nearly flat section. Thus even though escapes are rare, once beliefs do escape they are likely to move quite a distance (nearly 3 units) in the positive direction. This reflects what we observed in Figure 6, where escapes with $\rho = 0$ followed no definite pattern, while with $\rho = -1$ there were recurrent increases in the slope of roughly the same magnitude.

Figure 8 further illustrates the differences in escapes. Each panel of the figure fixes (r, ρ) and shows a histogram of the slope coefficient at the time of escape ($\gamma_1^\epsilon(\tau^\epsilon)$) from 5,000 simulated escapes for each of two different gain settings, $\epsilon = 0.1$ and $\epsilon = 0.01$. Also shown in each figure is the predicted escape point $\gamma_1(T)$ from the most probable escape path associated with that particular (r, ρ) setting. The top two panels show that for relatively small escapes, $r = \pm 0.5$, the distributions are relatively symmetric, with escapes in both the positive and negative directions. As the gain decreases, the escapes become more concentrated near the predicted points. The asymmetries we observed in Figure 7 are clear with the larger escape sets of $r = \pm 1.5$ and $r = \pm 3$ in the bottom panels of Figure 8. The figure only considers $\rho = -1$, as in the simulations we never (up to millions of periods) observed escapes of this magnitude for $\rho = 0$. For small gains, the escape distributions become highly concentrated around the predicted positive escape point. Thus the asymmetries in the rate functions are borne out in the simulations.

As above, the predictions of Theorem 4.1 are shown in Figure 9, which plots the escape distributions and escape times when $\rho = -1$ and $r = 3.1$ for varying gains, with 10,000 simulations for each gain setting. The left panels plot the predicted escape point for both the intercept (top panel) and slope coefficient (bottom panel), along with the median and bands covering 90% of the simulated distribution. As $\epsilon \rightarrow 0$ the




FIGURE 9. Escape distributions and escape times. The left panels plot the escape distribution of the coefficient $(\gamma_0^e(\tau^e), \gamma_1^e(\tau^e))$ for different gains (ε) , fixing $r = 3.1$ 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).

distribution tightens substantially around the prediction, and the median converges up to near the predicted level. The right panel of Figure 9 plots the log mean escape times from the simulations along with our prediction, and bands covering 90% of the simulated distribution. As in the first example, for small gains the predictions match the simulations quite well.

Finally, Figure 10 shows a summary of 10,000 simulated regular escape paths with $r = 3.1$, $\rho = -1$, $\varepsilon = \frac{1}{300}$, and $\mu_1 = 0.01$. The top panel plots the time paths of the intercept coefficient along an escape, while the bottom panel plots the slope. As above, the figure shows the predicted most probable path, and the paths closest to the mean, 5% and 95% quantiles of the simulated escape time distribution. All the escape paths are characterized by a relatively long period near the SCE, followed by a rapid increase in the slope and decrease in the intercept. As above, my results characterize the escape paths quite well, with the predicted most probable escape path being quite close to the mean from the simulations.

5.3. Interpretation and Relation to the Literature

The monopoly model exhibits strikingly different behavior depending on the value of ρ . When $\rho = 0$ escapes are symmetric, but 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_{1n} and W_{2n} are independent processes, there will occasionally be a sequence of correlated realizations. With $\rho < 0$,




FIGURE 10. Simulated and predicted regular escape paths with $r = 3.1$, $\rho = -1$, $\mu_1 = 0.01$ and $\varepsilon = \frac{1}{300}$ for the intercept coefficient γ_0 (top panel) and slope coefficient γ_1 (bottom panel). Each panel plots the path closest to the mean time (blue solid lines) and 90% band (black dashed lines) of the simulated distribution along with the predicted escape path (red dashed lines).

during such an episode output will have a smaller response to price changes than is typical. The firm interprets these outcomes as representing a decrease in the elasticity of its demand curve, and responds by raising its markup $a_n = h(\gamma)$, 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 learns the correct slope $\gamma_1 = b$ and sets the markup at the true profit maximizing level, and does not increase its markup any further. Eventually, there are sufficient uncorrelated 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, which in this case is equal to the SCE, and is 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 experience escapes when shocks push beliefs from the basin of attraction of one equilibrium to another. Notable

examples include evolutionary games, such as Kandori, Mailath, and Rob (1993) and Young (1993) and much subsequent literature, and macroeconomic models with multiple stable equilibria, such as Kasa (2004) and Ellison, Graham, and Vilmunen (2006). Williams (2014) uses the methods of this paper to study multiple equilibria. Similarly, the models of Marcet and Nicolini (2003) and Sargent, Williams, and Zha (2009) feature one stable equilibrium and an explosive region, with escapes occasionally causing beliefs to enter the explosive region.

More closely related to this paper are the learning models 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. Kolyuzhnov, Bogomolova, and Slobodyan (2014) also study this model using a continuous time approximation. Related models include Bullard and Cho (2005), McGough (2006), Ellison and Yates (2007), Sargent, Williams, and Zha (2006), Cho and Kasa (2008), several of which use the methods developed here. Similarly, Williams (2016) 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. Ellison and Scott (2013) build on this model to study volatility in oil prices. 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.

6. CONCLUSION

In this paper I 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. I have provided a precise characterization of these dynamics, and have illustrated how they can arise in examples. As the gain decreases to zero (across sequences), the beliefs converge in a weak sense to a self-confirming equilibrium. However ongoing stochastic shocks may occasionally lead beliefs to escape from the self-confirming equilibrium. I developed new theoretical methods to characterize the escape dynamics, and showed how to apply them in two simple economic models. I also showed how misspecification may generate locally self-reinforcing dynamics, which lead to recurrent large deviations from the self-confirming equilibrium.

The methods that I 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 Ellison, Graham, and Vilmunen (2006) and Cho and Kasa (2008)), fluctuations in prices resembling price wars (by Williams (2016)), and oil price volatility (by Ellison and Scott (2013)). Further, my results are not necessarily 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 (γ, R) into a the vector θ as in (11) and write (9)-(10) compactly as:

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

Then we define $\bar{\Psi}(\theta) = E\Psi(\theta, \xi_n)$ and $v_{n+1} = \Psi(\theta_n, \xi_{n+1}) - \bar{\Psi}(\theta_n)$. The following are the necessary assumptions for Theorem 3.1 above. We include an ε superscript on θ_n^ε and v_n^ε 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 , $\{\Psi(\theta_n^\varepsilon, \xi_{n+1})1_{\{\theta_n^\varepsilon \in A\}}; \varepsilon, n\}$ is *uniformly integrable*.²
3. The ODE $\dot{\theta} = \bar{\Psi}(\theta)$ has a point $\bar{\theta}$ which is *asymptotically stable*.³
4. The function $\bar{\Psi}(\theta)$ is *continuous*.
5. For each $\delta > 0$, there is a compact set A_δ such that $\inf_{n,\varepsilon} P(v_n^\varepsilon \in A_\delta) \geq 1 - \delta$.

Proof (Theorem 3.1). 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 $E\Psi(\theta_n^\varepsilon, \xi_{n+1}) = \bar{\Psi}(\theta_n^\varepsilon)$ is independent of ξ_{n+1} . This implies that the limit in (A8.1.9) is identically zero and that the β_n^ε 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 θ , $\Psi(\theta, \xi_{n+1})$ is a quadratic function of standard normal random variables. Therefore $P(|\Psi(\theta, \xi_{n+1})| \geq K) = f(K)$ 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 θ , we have that $P(|\theta_n^\varepsilon| \geq K) \rightarrow 0$ as $K \rightarrow \infty$, and so the sequence is tight. For part 2, note that $|\Psi(\theta_n, \xi_{n+1})|^2$ 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

This section collects assumptions and proofs for the results in Section 4. First we state the additional assumptions necessary for the large deviation principle.

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

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

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

$$\lim_{K \rightarrow \infty} \sup_n E(|X_n| 1_{\{|X_n| \geq K\}}) = 0.$$

³ A point \bar{x} is *asymptotically stable* for an ODE if any solution $x(t) \rightarrow \bar{x}$ as $t \rightarrow \infty$, 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 .

ASSUMPTIONS A.2.

1a. The sequence $\{|\Psi(\theta_n, \xi_{n+1})|\}$ is almost surely bounded by some constant $K < \infty$.

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

$$P(|\Psi(\theta_n, \xi_{n+1})| \geq s | \mathcal{F}_n) \leq B \exp(-s^\kappa) \text{ a.s.}$$

For Assumptions A.2, we require either part 1a or 1b to hold.

Proof (Theorem 4.1). (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.2, 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.2. 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 the theorem, the results follow 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). \blacksquare

A.3. CALCULATIONS FOR THE COST FUNCTIONS

A.3.1. Proof of Lemma 4.1

For use in the proof of Lemma 4.2 below, we develop the main calculations for a one-step version of H , then simplify to the static case. That is, we evaluate:

$$H_1(\theta, \alpha, y_0) = \log E_{y_0} \exp \langle \alpha, \Psi(\theta, \xi_1) \rangle$$

where $y_1 = \bar{A}y_0 + \bar{\Sigma}W_1$. For later use, we also assume here $W_1 \sim N(0, \Lambda)$, although in practice $\Lambda = I$. Using the definitions of the V matrices, we can write:

$$\begin{aligned} \exp(H_1) &= \exp\left(-\frac{1}{2}y_0'V_{yy}y_0\right)E_{y_0} \exp\left(-y_0'V_{yw}W_1 - \frac{1}{2}W_1'V_{ww}W_1\right) \\ &= \frac{1}{\sqrt{2\pi}}|\Lambda|^{-\frac{1}{2}} \exp\left(-\frac{1}{2}y_0'V_{yy}y_0\right) \int \exp\left(-y_0'V_{yw}W_1 - \frac{1}{2}W_1'V_{ww}W_1 - \frac{1}{2}W_1'\Lambda^{-1}W_1\right)dW_1. \end{aligned}$$

We now use a completing the square argument. We would like the terms inside the exponential to read $-\frac{1}{2}(W_1 - \mu_1)'\Omega_1^{-1}(W_1 - \mu_1)$ for some μ_1 and Ω_1 , in which case we would just be integrating a normal density and the integral would simply be unity. This holds if we set:

$$\Omega_1^{-1} = V_{ww} + \Lambda^{-1}, \quad \mu_1 = -\Omega_1 V_{yw}'y_0.$$

Making this substitution, adding and subtracting $-\frac{1}{2}\mu_1'\Omega_1^{-1}\mu_1$, and multiplying and dividing by $|\Omega_1|^{-\frac{1}{2}}$ gives:

$$\begin{aligned} \exp(H_1) &= \exp\left(-\frac{1}{2}y_0'V_{yy}y_0 + \frac{1}{2}\mu_1'\Omega_1^{-1}\mu_1\right)|\Omega_1|^{\frac{1}{2}}|\Lambda|^{-\frac{1}{2}} \cdot \\ &\quad \frac{1}{\sqrt{2\pi}}|\Omega_1|^{-\frac{1}{2}} \int \exp\left(-\frac{1}{2}(W_1 - \mu_1)'\Omega_1^{-1}(W_1 - \mu_1)\right)dW_1. \end{aligned}$$

The terms in the second line are thus the normal density, so we have:

$$H_1 = -\frac{1}{2}y_0'V_{yy}y_0 + \frac{1}{2}\mu_1'\Omega_1^{-1}\mu_1 + \frac{1}{2}\log|\Omega_1| - \frac{1}{2}\log|\Lambda| \tag{A.1}$$

$$= -\frac{1}{2}\log|\Lambda| - \frac{1}{2}\log|V_{yy} + \Lambda^{-1}| - \frac{1}{2}y_0'[V_{yy} - V_{yw}(V_{ww} + \Lambda^{-1})^{-1}V_{yw}']y_0. \tag{A.2}$$

Setting $\Lambda = I$ and $y_0 = 1$ gives the result (22).

A.3.2. Proof of Lemma 4.2

We follow the proof of Lemma 4.1 and proceed by induction. We already have determined the one-step version H_1 . Now we proceed by induction to determine the T step version H_T . That is, we suppose that H_T can be written:

$$H_T(\theta, \alpha, y_0) = \frac{1}{T} \log E_{y_0} \exp \left\langle \alpha, \sum_{n=1}^T \Psi(\theta, \xi_n) \right\rangle = \frac{1}{2T} \sum_{n=1}^T \log |\Omega_n| - \frac{1}{2T} y_0' \Theta_T y_0$$

for some sequences of matrices $\{\Omega_n, \Theta_n\}$. Note that this holds for H_1 with:

$$\Omega_1 = (V_{ww} + \Lambda^{-1})^{-1}, \quad \Theta_1 = V_{yy} - V_{yw}(V_{ww} + \Lambda^{-1})^{-1}V_{yw}'.$$

Now we evaluate H_{T+1} :

$$\begin{aligned} H_{T+1} &= \frac{1}{T+1} \log E_{y_0} \exp \left\langle \alpha, \sum_{n=1}^{T+1} \Psi(\theta, \xi_n) \right\rangle \\ &= \frac{1}{T+1} \log E_{y_0} \exp \langle \alpha, R^{-1} g(\gamma, \xi_1) \rangle E_{y_1} \exp \left\langle \alpha, \sum_{n=2}^{T+1} \Psi(\theta, \xi_n) \right\rangle, \end{aligned}$$

where the second line uses the law of iterated expectations. But notice that the second term is simply TH_T with the time indices simply shifted forward one period. Thus we can write:

$$\begin{aligned} \exp((T+1)H_{T+1}) &= \prod_{n=1}^T |\Omega_n|^{\frac{1}{2}} \exp(-\frac{1}{2} y_0' V_{yy} y_0) E_{y_0} \exp(-y_0' V_{yw} W_1 - \frac{1}{2} W_1' V_{ww} W_1 - \frac{1}{2} y_1' \Theta_T y_1) \\ &= \prod_{n=1}^T |\Omega_n|^{\frac{1}{2}} \exp(-\frac{1}{2} y_0' (V_{yy} + \bar{A}' \Theta_T \bar{A}) y_0) \cdot \\ &\quad E_{y_0} \exp(-y_0' (V_{yw} + \bar{A}' \Theta_T \bar{\Sigma}) W_1 - \frac{1}{2} W_1' (V_{ww} + \bar{\Sigma}' \Theta_1 \bar{\Sigma}) W_1) \end{aligned}$$

where in the first line we note that for the matrices the subscript indexes the number of steps in the iterations, not the date, and the second line uses the evolution of y_n . We then use the same complete the square argument as in the proof of Lemma 4.1, now setting:

$$\Omega_{T+1}^{-1} = V_{ww} + \bar{\Sigma}' \Theta_T \bar{\Sigma}, \quad \mu_{T+1} = -\Omega_{T+1} (V_{yw} + \bar{A}' \Theta_T \bar{\Sigma}).$$

Making these substitutions and doing the same manipulations as above, we now have:

$$\exp((T+1)H_{T+1}) = \prod_{n=1}^{T+1} |\Omega_n|^{\frac{1}{2}} \exp(-y_0' [V_{yy} + \bar{A}' \Theta_T \bar{A} - (V_{yw} + \bar{A}' \Theta_T \bar{\Sigma}) \Omega_{T+1} (V_{yw} + \bar{A}' \Theta_T \bar{\Sigma})'] y_0)$$

Thus we have:

$$H_{T+1} = \frac{1}{2(T+1)} \sum_{n=1}^{T+1} \log |\Omega_n| - \frac{1}{2T+1} y_0' \Theta_{T+1} y_0,$$

where Θ_n is determined recursively as:

$$\Theta_{n+1} = V_{yy} + \bar{A}' \Theta_n \bar{A} - (V_{yw} + \bar{A}' \Theta_n \bar{\Sigma}) \Omega_{n+1} (V_{yw} + \bar{A}' \Theta_n \bar{\Sigma})' \quad (\text{A.3})$$

with Ω_{n+1} a function of Θ_n as above. This proves the induction step and shows how to compute $\{\Omega_n, \Theta_n\}$.

To determine H we then take limits. Under the standard detectability and controllability conditions (see Kwakernaak and Sivan (1972)), Θ_n in (A.3) converges to a matrix Θ which satisfies the algebraic Riccati equation (24). This also implies Ω_n converges and so does the average in H_{T+1} . Noting as well that the effect of the initial condition y_0 dies out in the limit, we have the result.

A.4. VERIFYING THE ASSUMPTIONS IN THE MONOPOLY 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.2. Following our discussion after the proof of Theorem 3.1 above, we know that parts 1, 2, and 5 of Assumptions A.1 hold. Further, part 1b of Assumptions A.2 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 the text 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 γ . 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] Adam, Klaus, Albert Marcet and Juan Pablo Nicolini (2016) "Stock Market Volatility and Learning," *Journal of Finance*, 71: 33-82.
- [2] Benhabiv, Jess and Chetan Dave (2014) "Learning, Large Deviations and Rare Events," *Review of Economic Dynamics*, 17: 367-382.
- [3] Branch, William A. and George Evans (2011) "Learning about Risk and Return: A Simple Model of Bubbles and Crashes," *American Economic Journal: Macroeconomics*, 3: 159-191.
- [4] Bullard, James and In-Koo Cho (2005) "Escapist Policy Rules," *Journal of Economic Dynamics and Control*, 29: 1841-1865.
- [5] Cho, In-Koo and Kenneth Kasa (2008) "Learning Dynamics and Endogenous Currency Crises," *Macroeconomic Dynamics*, 12: 257-85.
- [6] Cho, In-Koo, Noah Williams, and Thomas J. Sargent (2002) "Escaping Nash Inflation," *Review of Economic Studies*, 69: 1-40.
- [7] Dembo, Amir and Ofer Zeitouni (1998) *Large Deviations Techniques and Applications*, Second Edition, Springer-Verlag, New York.
- [8] Dupuis, Paul and Harold J. Kushner (1987) "Stochastic Systems with Small Noise, Analysis and Simulation," *SIAM Journal on Applied Mathematics*, 47: 643-661.
- [9] 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.
- [10] Ellison, Martin, Liam Graham, and Jouko Vilmunen (2006) "Strong Contagion with Weak Spillovers," *Review of Economic Dynamics*, 9: 262-283.

- [11] Ellison, Martin and Andrew Scott (2013) “Learning and Price Volatility in Duopoly Models of Resource Depletion,” *Journal of Monetary Economics*, 60: 806-820.
- [12] Ellison, Martin and Tony Yates (2007) “Escaping Volatile Inflation,” *Journal of Money, Credit, and Banking*, 39: 981-993.
- [13] Evans, George and Seppo Honkapohja (2001) *Learning and Expectations in Macroeconomics*, Princeton University Press.
- [14] Evans, George and Seppo Honkapohja (2009) “Learning and Macroeconomics,” *Annual Review of Economics*, 1: 421-449.
- [15] Evans, George, Seppo Honkapohja, and Noah Williams (2010) “Generalized Stochastic Gradient Learning,” *International Economic Review*, 51: 237-262.
- [16] Fleming, Wendell and H. Mete Soner (1993) *Controlled Markov Processes and Viscosity Solutions*, Springer-Verlag, New York.
- [17] Freidlin, Mark I. and Alexander D. Wentzell (1998) *Random Perturbations of Dynamical Systems*, 2nd edition. Springer-Verlag, New York.
- [18] Fudenberg, Drew and David Levine (1998) *The Theory of Learning in Games*, MIT Press, Cambridge.
- [19] Kandori, Michihiro, George Mailath and Rafael Rob (1993) “Learning, Mutation, and Long Run Equilibria in Games,” *Econometrica*, 61: 29-56.
- [20] Kasa, Kenneth (2004) “Learning, Large Deviations, and Recurrent Currency Crises,” *International Economic Review*, 45: 141-173.
- [21] Kolyuzhnov, Dimitri, Anna Bogomolova, and Sergey Slobodyan (2014) “Escape Dynamics: A Continuous-Time Approximation,” *Journal of Economic Dynamics and Control*, 38: 161-183.
- [22] 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.
- [23] Kushner, Harold J. and George G. Yin (1997) *Stochastic Approximation Algorithms and Applications*, Springer-Verlag, New York.
- [24] Kwakernaak, Huibert and Raphael Sivan (1972) *Linear Optimal Control Systems*, John Wiley and Sons, New York.
- [25] Lucas, Robert E., Jr. (1978) “Asset Prices in an Exchange Economy,” *Econometrica*, 45: 1429-45.
- [26] Maier, Robert S. and Daniel L. Stein (1997) “Limiting Exit Distributions in the Stochastic Exit Problem,” *SIAM Journal on Applied Mathematics*, 57: 752-790.
- [27] Marcet, Albert and Juan Pablo Nicolini (2003) “Recurrent Hyperinflations and Learning,” *American Economic Review*, 93: 1476-1498.

- [28] 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.
- [29] McGough, Bruce (2006) "Shocking Escapes," *Economic Journal*, 116: 507-528.
- [30] Sargent, Thomas J. (1999) *The Conquest of American Inflation*, Princeton University Press.
- [31] Sargent, Thomas J. and Noah Williams (2005) "Impacts of Priors on Convergence and Escapes from Nash Inflation," *Review of Economic Dynamics*, 8: 360-391.
- [32] 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.
- [33] Sargent, Thomas J., Noah Williams, and Tao Zha (2009). "The Conquest of South American Inflation," *Journal of Political Economy*, 117: 211-256.
- [34] Timmermann, Alan G. (1993) "How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices," *Quarterly Journal of Economics*, 108: 1135-4563.
- [35] Williams, Noah (2014) "Long Run Equilibrium in Discounted Stochastic Fictitious Play," working paper, University of Wisconsin - Madison.
- [36] Williams, Noah (2016) "Equilibria, Escapes, and Price Wars," working paper, University of Wisconsin - Madison.
- [37] Worms, Julien (1999) "Moderate Deviations for Stable Markov Chains and Regression Models," *Electronic Journal of Probability*, 4(8): 1-28.
- [38] Young, H. Peyton (1993) "The Evolution of Conventions," *Econometrica*, 61: 57-84.
- [39] Young, H. Peyton (2004) *Strategic Learning and Its Limits*, Oxford University Press.