Sticky Thresholds Under Vigilant Learning

William A. Branch*
U.C. Irvine

September 28, 2026

Abstract

Many economic environments feature persistent regimes separated by unstable thresholds. This paper shows that model doubt can make these thresholds sticky. Under vigilant learning, agents forecast from a linear model while monitoring a richer diagnostic benchmark. Transitions can expose model misspecification that is nearly invisible within a regime. Rising concern leads agents to account for how uncertainty around their point forecasts interacts with a nonlinear economic environment – a Jensen adjustment. Because expectations affect outcomes, this response changes the economy's law of motion and can temporarily stabilize the threshold. Arrival at the threshold removes the predictability that originally raised concern, vigilance declines, and instability returns. Beliefs linger at the threshold, with a duration determined by the extent of vigilance. We characterize this three-phase sequence and illustrate the mechanism in growth and strategic-coordination environments. What appears to be an additional persistent regime may, in fact, be a temporarily stabilized boundary between regimes.

Keywords: thresholds, multiple equilibria, learning

1 Introduction

Many economic environments feature persistent regimes separated by a tipping threshold. Threshold-like behavior has been documented in technology adoption, poverty, neighborhood composition, and social coordination (Azariadis, 1996), together with closely related strategic complementarities in bank-runs and sovereign rollover crises (Vives, 2005). In

*Email: wbranch@uci.edu. I used ChatGPT (5.5 and 5.6 Sol) to help correct proofs and develop arguments used in Lemma 3; Claude Fable to generate simulation code; and Refine.ink for refinements that sharpened the main result. I take full responsibility for the paper's content.

canonical models a defining feature of these thresholds is instability with endogenous feedback carrying the state toward another regime once it crosses the threshold. Empirical tests of poverty traps, for instance, exploit the implication that the empirical distribution should feature little mass near a threshold or tipping point (Lybbert et al., 2004; Santos and Barrett, 2019; Balboni et al., 2022). This logic, though, assumes that the stability properties remain invariant during a transition between regimes. But, what happens if moving between regimes reveals misspecification in the approximating models agents use to forecast and take actions? This paper shows that model doubt can make an unstable threshold sticky.

There is a simple logic underlying the existence of sticky thresholds in dynamic, self-referential economies where actions depend on agents' expectations formed from an approximating model. Under standard point-expectations learning (Evans and Honkapohja (2001)), agents forecast from a parsimonious model and evaluate the nonlinear economic map at its point forecast. This certainty-equivalent approximation ignores how predictive uncertainty around the point forecast interacts with the map's curvature. A vigilant learner, on the other hand, still forecasts with that approximating model while also monitoring a richer model that serves as a diagnostic benchmark. The benchmark is not a candidate to replace the forecasting model. Its role is to assess when the forecasting model omits predictable variation and how dogmatic the agent should be in their point forecast. Following Lanzani (2025), concern about misspecification is endogenous, with vigilance rising as the benchmark outpredicts it and declining otherwise. When misspecification concern rises, agents partially restore the non-linear curvature effect omitted by point expectations through a second-order Jensen adjustment. That correction combines the local curvature of the map with a Jensen intensity that weights the forecast-error variance by vigilance and the agent's sensitivity to misspecification. Because the model is self-referential, this vigilance changes the actual law of motion and can reverse the threshold's stability.

This mechanism is particularly transparent in an S-shaped environment: see Figure 1. In the symmetric baseline case, the middle fixed point coincides with the inflection point and repels beliefs under certainty-equivalent learning. The Jensen adjustment increases expected outcomes below the inflection point and lowers them from above, creating a force towards the threshold from both sides. When vigilance towards misspecification is low, there are two stable outer regimes and the middle unstable threshold. When the Jensity intensity rises above a critical value, the outer regimes collapse and the threshold becomes attracting. Once the economy arrives at the threshold, however, the transitional predictability that elevated concern in the first place disappears, the benchmark loses its statistical advantage, vigilance declines, the outer regimes reappear, and the threshold is again unstable. Beliefs are released from the sticky threshold, but only after a stretch of time whose duration depends on the cumulative excess Jensen intensity that pinned them toward the threshold, that is, the episode's vigilance budget.

The main theoretical results formalize this mechanism in a scalar self-referential economy in which current outcomes depend on expectations of a non-linear function of future outcomes. Agents estimate a simple linear perceived law of motion in real time using, for instance, discounted least squares. Under certainty-equivalent learning, the temporary-equilibrium map has two expectationally stable ("E-stable") outer fixed points separated by an E-unstable middle threshold.[1] See Figure 1 for an illustration. Constant-gain learning, then, spends long periods near an outer regime, interrupted by occasional unlikely sequences of shocks that drive beliefs toward or across the separatrix. Vigilant agents also forecast from that perceived linear model but compare it continuously against a richer diagnostic benchmark. When the benchmark outpredicts the PLM, endogenous concern increases and agents partially relax their certainty-equivalence by applying a Jensen adjustment to the nonlinear function.

The first theoretical result establishes a clear distinction between the data generated within a regime and during a transition towards the threshold. Near either outer equilibrium, the benchmark's advantage is due to constant-gain learning noise, and vigilance fades with the gain parameter. Thus, the two E-stable equilibrium landscape persists with probability approaching one for small learning gains. Along a gradual transition path, however, the conditional mean of the state is moving and creating predictable comovement that the benchmark is able to detect. Importantly, the accumulated positive log-likelihood ratio is bounded away from zero even for small gains, implying that the accumulated evidence does not dissipate. So, conditional on a transition, a sufficiently high sensitivity to misspecification will cause the Jensen intensity to cross the pitchfork threshold illustrated in Figure 1 with probability approaching one, thereby, eliminating the outer equilibria and making the threshold attracting.

[1] Evans and Honkapohja (2001) show that the stability under real-time learning for a broad class of learning algorithms is governed by a simple and intuitive condition known as E-stability. E-stability is the local stability condition for a resting point of an ordinary differential equation (ODE) that describes how coefficients update based on the gap between the true population coefficient and the agent's estimate.

Then the second result shows that, conditional on entering that vigilance region, high vigilance continues to attract beliefs toward the threshold. Key to the cycle is that once the system arrives at that region the signal that initially raised agents' misspecification concerns vanishes. The approximating model and the benchmark become indistinguishable here, leading to declining vigilance until, eventually, the outer equilibria reemerge. Beliefs, though, linger at the threshold. High vigilance holds these beliefs in place, and once it declines the threshold becomes unstable again, though it takes time for beliefs to be pushed away. The duration of a lingering episode depends on the vigilance budget, how much excess vigilance accumulates above the threshold. Eventually, lingering ends and the economy returns to one of the outer equilibria. A third result constructs a finite-action large-deviation path, as in Williams (2026), that runs from an outer regime to the collapse regime and, under certain conditions, sticky-threshold episodes are recurrent.

For tractability, the formal results are developed with a symmetric cubic map. However, all the lingering mechanism needs is a locally unstable fixed point at an inflection, curvature in the map directing the Jensen adjustment inward, and then enough endogenous vigilance to reverse the local stability properties. The paper develops additional results in the context of a growth economy with multiple balanced growth paths and a strategic-coordination game. With serially correlated noise, it is possible for nonlinear misspecification to support high vigilance as an equilibrium phenomenon and not just during a transition. Moreover, in simulations with finite gains $\gamma > 0$ and sufficient sensitivity to misspecification, learning noise can keep vigilance high at a sticky threshold for even longer stretches of time.

Related literature. The theory and results here are related to several literatures on misspecified beliefs and learning. Marcet and Sargent (1989) construct an equilibrium concept, in the spirit of rational expectations, for an environment in which some state variables are hidden from agents. Restricted perceptions equilibria (RPE) are surveyed in Evans and Honkapohja (2001), Branch (2006), and Branch and McGough (2018). In these approaches, agents form expectations from parsimonious perceived laws of motion that may omit lags, state variables, nonlinearities, or other features of the true stochastic process, which is itself determined in part by those expectations. Although the forecasting model is misspecified, at an RPE its coefficients are optimal within the restricted class. RPEs, therefore, preserve the cross-equation restrictions that are a salient feature of rational ex-

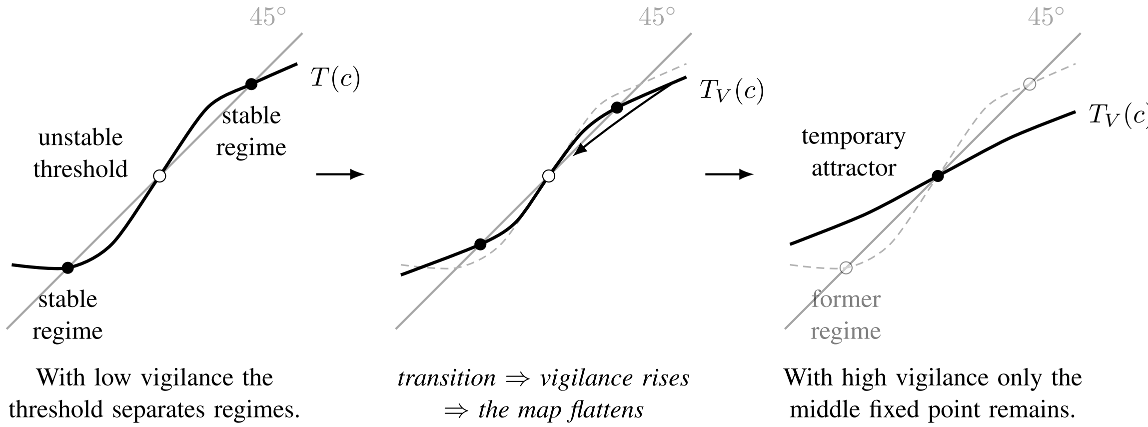

Figure 1: Sticky thresholds in the temporary-equilibrium mapping from beliefs (c). The gray line is the $45°$ line and the black curve is the temporary-equilibrium map. Low vigilance: two stable outer fixed points and an unstable threshold. Vigilance rising flattens the map. Enough vigilance, middle threshold is a temporary attractor.

pectations while allowing agents to utilize underparameterized models.[2]

An RPE is also closely related to the Berk–Nash equilibrium concept of Esponda and Pouzo (2016), Esponda et al. (2021), and Fudenberg et al. (2021). In both approaches, beliefs are disciplined by requiring the agents' model to be best-fitting within a restricted subjective class under data generated endogenously by those beliefs and actions. The contribution here is to make confidence in that restricted class an endogenous state variable. A forecasting model may fit well within each persistent regime but fail along the paths between them. When transition data reveal that the restricted class is omitting predictable variation, the resulting concern changes how agents use their model. In a self-referential economy that response changes the actual law of motion itself and can alter the number and stability of equilibria.[3]

The motivation for endogenous concern comes from Lanzani (2025), who studies an agent who entertains a set of structured models but worries that the true data-generating process may lie outside that restricted class. Vigilance rises and falls with statistical ev-

[2]Other restricted-perceptions approaches are developed in Sargent (1999), Hommes and Zhu (2014), Hommes and Sorger (1997), and Branch et al. (2022).

[3]The RPE used here is also closely related to the stochastic consistent expectations equilibrium of Branch and McGough (2005), Hommes et al. (2013), and Bullard et al. (2008). For the purely autoregressive PLMs considered in this paper, the OLS normal equations coincide with the corresponding Yule–Walker moment conditions (through lag p), so the two equilibrium concepts coincide. We therefore use the RPE terminology standard in the adaptive-learning literature.

idence against the maintained models, measured by a likelihood ratio relative to a richer class of alternatives. This provides a disciplined account of when agents should become less confident in their preferred specification. We adopt that discipline in a different environment and for a different purpose. Here the structured model is the agent's parsimonious perceived law of motion, the richer model is an auxiliary diagnostic benchmark, and vigilance determines the size of a Jensen adjustment for uncertainty surrounding the point forecast. In Lanzani (2025), endogenous concern changes decisions within a given environment. Here beliefs feedback into outcomes, outcomes determine the evidence against the perceived law of motion, and then that evidence determines the vigilance that enters beliefs. Through this feedback, concern about misspecification can change the equilibrium landscape that generates the evidence in the first place.

The Jensen adjustment also fits previous models where risk/uncertainty adjusts point expectations. A risk-adjustment is optimal under an asymmetric loss and can explain survey biases (Capistran and Timmermann, 2009). In asset-pricing applications, (Branch and Evans, 2011; Williams, forthcoming) let agents estimate risk alongside return and generate asset-price bubbles/crashes or Minsky-style cycles. Here the second-moment arises through vigilance to misspecification and an evidence-driven measure of misspecification concern.

The model-comparison device is also closely related to the model-validation framework of Cho and Kasa (2015). While model adjustment is discrete in Cho–Kasa and continuous here, the important distinction is about model-selection (model validation) versus model confidence within a class (vigilance). When an agent's forecasting model in Cho and Kasa (2015) fails a specification test, the agent selects a different model from that class, and use the new model to form expectations. The benchmark here is not a replacement to the agent's model, instead, the benchmark's likelihood advantage accumulates into vigilance, which, along with forecast-error variance and concern sensitivity, determines the size of the Jensen adjustment.[4]

Finally, the formalization of the three-phase sticky-threshold dynamics builds on the small-gain escape analysis of (Williams, 2019, 2026; Kasa, 2004). The most closely re-

[4]Another related finding is Cho and Libgober (2025) who endogenize the choice of a parsimonious misspecified representation, trading off decision payoffs against model complexity. Their object is the design of the parsimonious model and whether it is underparameterized. Here we take the restricted forecasting class as given and study evidence-driven variation in confidence in that class.

lated dynamic feature is the *fold remnant* in Williams (2026). After an exogenous parameter change destroys a stable–unstable pair of equilibria, the mean dynamics can retain a non-equilibrium point at which the drift is weak. The escape-cost function is nearly flat there and, under the conditions in Williams (2026), escapes from the surviving equilibrium concentrate near the remnant. A sticky threshold differs in three respects. First, the source of the geometry is different. A fold remnant is generated by an exogenous change in the drift, while a sticky threshold is an ordinarily repelling fixed point whose local stability is reversed endogenously by the agent's own vigilance. Second, the underlying mechanism governing persistence is different. The behavior near a fold remnant is governed by the exogenously parameterized drift and escape-cost landscape, whereas lingering at a sticky threshold is governed by an endogenous vigilance budget accumulated during the episode. Third, both mechanisms pose related identification problems. Williams (2026) shows that data generated by multiple stable equilibria can resemble data from a unique-equilibrium model with a fold remnant. Here, what appears to be a third regime may instead be a temporarily stabilized separatrix between two persistent regimes.

The paper proceeds as follows. Section 2 defines vigilant learning. Section 3 specializes to the symmetric cubic and shows sticky thresholds via stochastic simulation. Section 4 provide the analytic results. Section 5 presents extensions and applications.

2 A self-referential learning framework with vigilance

This section presents the framework abstractly. We describe the class of self-referential models, define a Restricted Perceptions Equilibrium (RPE), introduce the vigilance mechanism, and define an RPE with vigilance. The framework is deliberately general, while the next section focuses on a tractable map with i.i.d. shocks and an intercept-only PLM.

2.1 Self-referential models and the linear PLM

Consider a scalar economy

$$x_t = \hat{\mathbb{E}}_t G(x_{t+1}) + \varepsilon_t, \tag{1}$$

where x_t is an endogenous state, $G : \mathbb{R} \to \mathbb{R}$ is a nonlinear structural map, and ε_t is an exogenous stationary shock process. We assume G is at least twice continuously differentiable with bounded derivatives on compact sets. The conditional expectation $\hat{\mathbb{E}}_t$ is taken with respect to the agent's information at t and reflects (possibly) non-rational expectations.[5]

A standard simplification combines point expectations with a linear perceived law of motion. Point expectations replace the conditional expectation of a nonlinear function with the function evaluated at the point forecast,

$$\hat{\mathbb{E}}_t G(x_{t+1}) \approx G(x^e_{t+1}), \qquad x^e_{t+1} = \hat{\mathbb{E}}_t x_{t+1}. \tag{2}$$

Equation (2) is exact if G is linear or the conditional distribution of x_{t+1} is degenerate.

Agents form x^e_{t+1} from a linear AR(p) perceived law of motion (PLM), $x_{t+1} = c + b(L)x_t + \nu_{t+1}$, estimated in real time via constant-gain least squares. Letting $\theta = (c, b(L))'$ collect the PLM coefficients, define $z_t = (1, x_{t-1}, \ldots, x_{t-p})'$ for $p \geq 1$ and $z_t = 1$ for $p = 0$. The constant-gain recursion is

$$\theta_{t+1} = \theta_t + \gamma R^{-1}_{t+1} z_t(x_t - \theta'_t z_t), \qquad R_{t+1} = R_t + \gamma(z_t z'_t - R_t)\,, \tag{3}$$

with gain $0 < \gamma < 1$. This fits the standard stochastic-approximation framework studied by Evans and Honkapohja (2001). After plugging point expectations (2) into the structural equation (1), the recursion (3) defines an actual law of motion (ALM) for x_t.[6]

2.2 Restricted perceptions equilibrium

The linear PLM is, in general, misspecified — it cannot match a nonlinear ALM at every horizon — but agents pick the best linear forecast within their restricted regressor set, thus preserving the salience of cross-equation restrictions. That forecast model satisfies the least-squares orthogonality condition $\mathbb{E} z_t(x_t - \theta' z_t) = 0$, which implicitly defines a map $T : \mathbb{R}^{p+1} \to \mathbb{R}^{p+1}$ taking PLM coefficients to the projection coefficients defined by the ALM arising from those PLM coefficients.

Definition 1 (Restricted Perceptions Equilibrium)**.** A vector $\theta^* \in \mathbb{R}^{p+1}$ is a *Restricted Perceptions Equilibrium* (RPE) if it is a fixed point of the T-map: $\theta^* = T(\theta^*)$.

[5] The framework also accommodates noise inside the structural map, $G(x_{t+1}, \varepsilon_t)$, with appropriate modifications to the equations below.

[6] A timing convention is that the time-t information includes endogenous variables up to $t-1$ and exogenous shocks up to t. For $p \geq 1$, x^e_{t+1} is the PLM forecast iterated two steps ahead given predetermined z_t. The intercept-only baseline is unaffected by this convention.

At an RPE, the PLM is the best linear predictor of x_t given the restricted regressor set, under the ALM induced by the PLM itself. RPE generalizes rational expectations to settings where agents' models are linear restrictions on a nonlinear environment (Branch and McGough (2005), Hommes et al. (2013)).

2.3 The vigilance mechanism

The misspecification intrinsic to (2) with the linear PLM has two components. The linear PLM when the data are generated from a nonlinear ALM. And, importantly, imposing point-expectations ignores the curvature of G in expectation. Following Lanzani (2025), we model an agent who is statistically sophisticated and detects evidence against their PLM and partially corrects for the curvature they are missing. There are three components to the vigilance mechanism: correction, detection, and the level of vigilance towards misspecification.

Jensen adjustment. When evidence accumulates against the PLM, the agent replaces the certainty-equivalent forecast (2) with a second-order Taylor approximation:

$$\hat{\mathbb{E}}_t G(x_{t+1}) \approx G(x^e_{t+1}) + \tfrac{1}{2} G''(x^e_{t+1}) V_t, \tag{4}$$

where $V_t \geq 0$ is a *Jensen intensity* that scales the second-order correction. We parameterize the intensity as

$$V_t = \kappa \, \Lambda_{t-1} \, \widehat{\sigma}^2_{t-1}, \tag{5}$$

where $\Lambda_{t-1} \geq 0$ is the agent's vigilance state, $\widehat{\sigma}^2_{t-1}$ is the PLM forecast variance estimated from data through $t-1$, and $\kappa \geq 0$ is the concern sensitivity.[7] When $\Lambda_{t-1} = 0$ it reduces to point expectations and when $\Lambda_{t-1} > 0$ the agent incorporates a misspecification concern-weighted spread around the PLM-implied conditional mean.

The Jensen intensity is not an estimate of the innovation variance associated with the agent's PLM. Rather, it is the variance of a misspecification concern-based perturbation to the conditional mean implied by the PLM. To see this behavioral interpretation, let m_t

[7] This particular form assumes the Jensen adjustment is linear in the vigilance state. An alternative (bounded) form is $\tilde{V}_t = h(V_t) \equiv \bar{V} \frac{V_t}{\bar{V}+V_t}$. Since $h(V) = V - \frac{V^2}{\bar{V}} + O(V^3)$, the key Phase 1 results below do not hinge on the linear form assumed here. All that is required is that $h(0) = 0, h'(0) = 1$, and $\bar{V}$ big enough.

denote the agent's PLM-derived point forecast and suppose that, when evaluating the nonlinear map, the agent entertains a symmetric perturbation u_{t+1} to that conditional mean satisfying $\widehat{E}_t u_{t+1} = 0$, $\widehat{E}_t u_{t+1}^2 = V_t$. The agent then evaluates $\widehat{E}_t G(m_t + u_{t+1}) \approx G(m_t) + \frac{1}{2} G''(m_t) V_t$, using the second-order approximation for a smooth G. For a cubic G, symmetry of the perturbation makes this expression exact.[8] This gives a behavioral interpretation to the Jensen adjustment. The benchmark is a diagnostic which reveals that forecasting class is misspecified, but it is not viewed as credible enough to replace the PLM. Instead, any likelihood advantages accumulate over time. Vigilance expands the range of deviations from the PLM-implied conditional mean and the agent then evaluates G under that mean-preserving spread. Under this interpretation, $\widehat{\sigma}_{t-1}^2$ provides the scale of recent PLM forecast errors, Λ_{t-1} measures accumulated evidence that the restricted class is misspecified, and κ translates that evidence into the Jensen adjustment.

Detection. The agent measures evidence against the PLM by maintaining a richer nested predictive model alongside it. The added flexibility may come from extra lags, nonlinear transformations, interactions, or other low-dimensional forecasting terms. The canonical baseline example is an AR($p+1$) benchmark for an AR(p) PLM. Both models are estimated in real time with the same constant gain γ.[9] At each t, the agent computes the Gaussian log-likelihood ratio (LLR)

$$\ell_t = \log \frac{f_{\text{bench}}(x_t \mid \mathcal{F}_{t-1}; \widehat{\vartheta}_t)}{f_{\text{plm}}(x_t \mid \mathcal{F}_{t-1}; \widehat{\theta}_t)}, \tag{6}$$

where $\widehat{\theta}_t$ and $\widehat{\vartheta}_t$ collect the recursively estimated PLM and benchmark parameters. When $\ell_t > 0$, the benchmark's predictive likelihood is higher, and we say it outpredicts the PLM in period t.

[8] For example, let $u_{t+1} = \pm\sqrt{V_t}$ with equal probability. Then $\widehat{E}_t G(m_t + u_{t+1}) = \frac{1}{2}\left\{G(m_t + \sqrt{V_t}) + G(m_t - \sqrt{V_t})\right\}$, which equals $G(m_t) + \frac{1}{2} G''(m_t) V_t$ for any cubic G.

[9] Equal gains is a substantive modeling choice that reflects a convention that agents update their models at a consistent rate. A benchmark updated on a faster time-scale than the PLM would simplify parts of the analysis as the benchmark would then better track its population coefficients. However, we do not follow this approach in order to remain closer to the large-deviation principle in Williams (2019).

Vigilance accumulation. Vigilance accumulates this evidence via a one-sided exponentially weighted moving average:

$$\Lambda_t = (1-\gamma)\Lambda_{t-1} + \gamma \max(\ell_t, 0). \tag{7}$$

Vigilance moves toward ℓ_t^+. In line with Lanzani (2025), evidence against the PLM raises concern, so accumulated vigilance measures how often the benchmark outpredicts the PLM.

2.4 RPE with vigilance

Vigilance changes the actual law of motion since the vigilant forecast (4) replaces $G(x_{t+1}^e)$ in (1). Define the vigilance-adjusted T-map as follows. For any $V \geq 0$, $T_V : \mathbb{R}^{p+1} \to \mathbb{R}^{p+1}$ takes PLM coefficients to the projection coefficients implied by the ALM $x_t = G(x_{t+1}^e) + \frac{1}{2}G''(x_{t+1}^e)V + \varepsilon_t$. At $V = 0$, $T_V = T$ recovers the standard T-map. At $V > 0$, the curvature term alters the ALM and hence the projected T-map. In equilibrium, the vigilance-augmented system pins down the PLM coefficients, the benchmark coefficients, and the vigilance state simultaneously.

Definition 2 (RPE with vigilance)**.** A tuple $(\theta^*, \vartheta^*, \Lambda^*)$ is an *RPE with vigilance* if

(i) $\theta^* = T_{V_{\mathrm{RPE}}^*}(\theta^*)$, where $V_{\mathrm{RPE}}^* = \kappa\Lambda^*\widehat{\sigma}^2(\theta^*)$ and $\widehat{\sigma}^2(\theta^*)$ is the stationary PLM forecast variance under the ALM induced by θ^* and V_{RPE}^*;

(ii) ϑ^* is the population-optimal benchmark coefficient vector under the same ALM;

(iii) $\Lambda^* = \mathbb{E}[\max(\ell_t^*, 0)]$, where ℓ_t^* is the LLR (6) evaluated at (θ^*, ϑ^*) under the vigilance-adjusted ALM.

Vigilance Λ^* depends on the stationary LLR, the LLR depends on the PLM and benchmark coefficients, those coefficients depend on the ALM, and the ALM depends in turn on Λ^* through the Jensen intensity V_{RPE}^*. When $\kappa = 0$, outcomes and beliefs collapse to Definition 1. When $\kappa > 0$, vigilance can shift the location, stability, and even the existence of fixed points relative to the standard RPE. This is exactly the mechanism studied in the remainder of the paper. When the maintained PLM is correctly specified at the fixed point, the benchmark gains nothing, the stationary LLR is zero, and Definition 2 reduces to Definition 1 with $\Lambda^* = 0$.[10] Nontrivial population vigilance instead requires misspecification that leaves the benchmark with a strict predictive advantage under the stationary ALM.[11]

3 Model

This section introduces vigilance learning into a non-linear self-referential univariate model.

3.1 Economic law of motion

We now adopt the scalar framework studied in Evans and Honkapohja (2001, Ch. 11), with point expectations (2) and a linear PLM, making three simplifying assumptions that deliver tractable analytic results.

First, assume that $\varepsilon_t \sim i.i.d.(0, \sigma^2)$. This implies that there exist equilibria that are "noisy steady-states" $x_t = \bar{x} + \varepsilon_t$, $\bar{x} = \mathbb{E}G(\bar{x} + \varepsilon_t)$. Evans and Honkapohja (2001) establish the existence of noisy steady-state equilibria for a small noise limit. Second, the possibility of noisy steady-state equilibria motivates the linear perceived law of motion: $x_t = c + \varepsilon_t$. This is a restricted perceptions model in the simplest possible sense. Section 5 illustrates what changes with asymmetric $G(\cdot)$, persistent shocks, and richer forecasting models. Agents update their intercept using constant-gain least-squares,

$$c_{t+1} = c_t + \gamma\,(x_t - c_t) \tag{8}$$

with gain $0 < \gamma < 1$.

Third, we assume a cubic form for the map G in (1): $G(x) = x - \frac{(x-\alpha_L)(x-\alpha_M)(x-\alpha_H)}{k}$, with $\alpha_L < \alpha_M < \alpha_H$. The analytic results exploit symmetry when $\alpha_M = {}^{(\alpha_L+\alpha_H)}/_2$, and $k > 0$ chosen so that $|G'(\alpha_L)|, |G'(\alpha_H)| < 1$ while $G'(\alpha_M) > 1$. Symmetry has the convenient property that $G''(\alpha_M) = -\frac{2}{k}\,(2\alpha_M - \alpha_H - \alpha_L) = 0$. Define $\Delta = \alpha_H - \alpha_L$.

[10] This is the case in the cubic example below where the RPEs are effectively noisy steady-states.

[11] Definition 2 is a population-level object. Section 5 illustrates another way to sustain positive vigilance. With a fixed positive gain parameter, one-sided LLR updating can sustain positive vigilance under the learning limiting distribution even when the equilibrium advantage is zero. We call this a finite-gain vigilance regime, not an RPE with vigilance.

3.2 A micro-foundation: platform monopoly with network externalities

The cubic map G used in the analysis can be given a simple micro-foundation using a platform monopolist facing consumers with forward-looking, nonlinear network externalities. This subsection details a simple environment inspired by Katz and Shapiro (1985) and Economides and Himmelberg (1995).

A platform monopolist sets a per-period fee p_t. A unit mass of consumers have heterogeneous intrinsic values $\theta_i \sim U[0,1]$. Consumer i will join the platform whenever their total surplus exceeds the fee, $\theta_i + \hat{\mathbb{E}}_t v(x_{t+1}) \geq p_t$, where x_{t+1} is the future network size and $v(\cdot)$ is the network value function that captures the spillover benefit from other users. Consumers here are forward-looking in the sense of Katz and Shapiro (1985) where current platform adoption depends on expectations about the future spillover benefit. This is what creates the self-referentiality.

Since every consumer with $\theta_i \geq p - \hat{\mathbb{E}}_t[v(x_{t+1})]$ joins the platform, aggregate demand at fee p is $x = 1 - p + \hat{\mathbb{E}}_t[v(x_{t+1})]$. A monopolist, who takes consumer expectations as given, maximizes $\pi = px$, yielding $p^* = \left(1 + \hat{\mathbb{E}}_t[v(x_{t+1})]\right)/2$. It follows that demand is $x_t = \left\{1 + \hat{\mathbb{E}}_t[v(x_{t+1})]\right\}/2$. Defining $G(x) \equiv (1 + v(x))/2$, this is $x_t = \hat{\mathbb{E}}_t[G(x_{t+1})]$, which is the self-referential form (1). The agent's forecasting problem reduces to evaluating the expectation of the nonlinear map G. The nonlinearity inherited from the network value function v is precisely what makes this evaluation nontrivial.

Now we parameterize the platform spillover benefit to yield a cubic map. In any steady state $v(x^*) = 2x^* - 1$. To generate three steady states at, say, $\alpha_L, \alpha_M, \alpha_H$, one possibility is $v(x) = (2x-1) - A(x-\alpha_L)(x-\alpha_M)(x-\alpha_H)$, which yields the cubic G-map with $k = {}^2\!/_A$. The network value $v(x)$ combines a linear network effect $2x - 1$ with a cubic term that captures threshold effects and congestion. There is a positive marginal network value, $v'(x) \geq 0$ on $[\alpha_L, \alpha_H]$ under symmetry, provided that $k > \Delta^2/2$, a restriction that holds in the paper's calibration.[12] At the middle steady state, the spillover benefits just offset the stand-alone loss, while at the high steady state the platform generates strong spillover network value.

[12]The participation rate is also constrained to $[0,1]$, on the domain with bounded expected returns. Outside of that, demand is censored, the T-map flattens, which just reinforces the outer wells.

3.3 Expectational stability

To assess the stability of the restricted perceptions equilibria, we now impose point expectations so that the actual law of motion is $x_t = G(c_t) + \sigma\eta_t$, and beliefs follow the constant-gain recursion (8), which is simply an estimate of the mean value of x_t generated with geometrically declining weights. The combination of point expectations, and a PLM consisting of just an intercept term, is called "steady state learning." Combining the ALM and (8) leads to $c_{t+1} = (1-\gamma)c_t + \gamma G(c_t) + \gamma\sigma\eta_t$ with $\eta_t \sim N(0,1)$. Evidently, local stability of this recursion near $c^* = G(c^*)$ requires $|1-\gamma+\gamma G'(c^*)| < 1$ (or $G'(c^*) < 1$ as $\gamma \to 0$). But, we are interested in a more general stability concept that applies to other stochastic approximation algorithms.

Evans and Honkapohja (2001) define Expectational Stability ("E-stability") as local stability of the ordinary differential equation $\dot{c} = G(c) - c$. The interpretation is as follows. Suppose beliefs are fixed at $c_t = c$ and a long history of data on x_t is generated. The sample mean converges to $G(c)$, so any reasonable learning algorithm adjusts c in the direction of $G(c)-c$. If this adjustment contracts toward a fixed point, that RPE is E-stable. Linearizing this ODE, the condition is $G'(c^*) < 1$, which implies that (1.) the outer fixed points α_L, α_H are E-stable; and, (2.), the middle fixed point α_M is E-unstable. For the intercept-only PLM, the E-stability condition coincides with the difference-equation condition derived above.[13]

3.4 Vigilance learning: relaxing certainty equivalence

Standard point-expectations learning evaluates the nonlinear map G at a point forecast. In a nonlinear environment, however, evaluating G at a point forecast ignores the effect of forecast uncertainty on expected outcomes. Since agents here use a linear perceived law of motion in a nonlinear environment, their model is misspecified.

This motivates the return to vigilant learning. Our departure from the standard approach is deliberately modest. Agents continue to use a linear PLM, but when the data suggest that this specification fits poorly, they partially account for the uncertainty surrounding their

[13] In this paper, there is one particular advantage coming from the ODE formulation. When E-stable, Evans and Honkapohja (2001) show that the continuous-time interpolation of the learning dynamics (for a range of learning algorithms) tracks solutions of the mean-dynamics ODE on the slow (order $1/\gamma$) time-scale. This makes the mean-dynamics ODE the natural starting point for the large deviations analysis in Section 4, which characterizes the global dynamics of an escape from the basin of attraction of an E-stable RPE.

forecast. The Jensen adjustment (4)–(5) applied here with $x^e_{t+1} = c_t$, the actual law of motion becomes

$$x_t = G(c_t) + \frac{1}{2}G''(c_t)V_t + \sigma\eta_t \equiv T_{V_t}(c_t) + \sigma\eta_t, \tag{9}$$

where $V_t = \kappa\Lambda_{t-1}\widehat{\sigma}^2_{t-1}$ is the effective Jensen intensity. Both factors are functions of data through $t-1$, so V_t is predetermined when x_t is realized. The term $\frac{1}{2}G''(c_t)V_t$ captures, to second order, the effect of integrating the nonlinear map G over forecast uncertainty rather than evaluating it at a point forecast. The cubic specification is especially convenient here. Because agents share a common linear forecasting model, their subjective beliefs are centered at the linear forecast, and under Gaussian forecast errors the second-order expression is exact for cubic G. The concern sensitivity $\kappa \geq 0$ parameterizes the strength of the vigilant reaction, and when $\Lambda_{t-1} = 0$ the standard ALM emerges.

Beliefs here are misspecified in two senses. First, the linear PLM is correct on, but not off, the equilibrium path. Each E-stable RPE (Definition 1) is a noisy steady state with the intercept-only PLM coinciding with the ALM, a self-confirming equilibrium. Away from it, the nonlinearity moves the conditional mean and the linear PLM is misspecified. When the gain γ and the shock variance σ^2 are small, beliefs concentrate near a single fixed point and this off-path misspecification is a second-order concern. Concern being small near an E-stable RPE is what the escape dynamics exploit and why the economy stays near one fixed point for long stretches.

The second form of misspecification comes from the constant-gain learning. Under the intercept-only PLM agents believe x_t is i.i.d. around a constant. Under learning, however, constant-gain updating makes the intercept c_t drift, which introduces persistence into the actual law of motion that the PLM cannot capture. Linearized around a stable RPE, x_t is ARMA(1,1) with first-order autocorrelation proportional to γ (see Section 4). Even though the autocorrelation is small an agent comparing the fit of the i.i.d. PLM against an AR(1) alternative will find persistent evidence against the PLM model. This is the econometric signal that Λ_t is designed to detect.[14]

[14] A natural question is whether the AR(1) benchmark is the right comparison model. For the intercept-only PLM the dominant detectable misspecification is the autocorrelation (proportional to γ) from the hidden drifting state, and an AR(1) benchmark is its natural estimator, while the nonlinear channel is second-order when σ and γ are small. Richer PLMs call for correspondingly richer benchmarks, which we explore in the applications.

3.5 Misspecification detection and vigilance accumulation

In the simple example of a constant-only PLM an AR(1) is the natural benchmark comparison model

$$x_t = a_{c,t} + b_{c,t}\, x_{t-1} + \varepsilon_t^c. \tag{10}$$

As before, the PLM recursion is (8). Letting $\theta_{c,t} = (a_{c,t},\, b_{c,t})'$ and $z_{c,t} = (1,\, x_{t-1})'$, the comparison model's coefficients are updated via constant-gain recursive least squares:

$$R_{c,t+1} = R_{c,t} + \gamma\big(z_{c,t}\, z_{c,t}' - R_{c,t}\big)\,, \tag{11}$$

$$\theta_{c,t+1} = \theta_{c,t} + \gamma\, R_{c,t+1}^{-1}\, z_{c,t}\, e_t^c, \tag{12}$$

where $e_t^c = x_t - \theta_{c,t}' z_{c,t}$ is the comparison model's forecast error and $R_{c,t}$ is the sample second moment matrix. Both the PLM and the benchmark model also update variance estimates with the same gain:

$$\widehat{\sigma}_t^2 = \widehat{\sigma}_{t-1}^2 + \gamma\big(e_t^2 - \widehat{\sigma}_{t-1}^2\big)\,, \tag{13}$$

$$\widehat{\sigma}_{c,t}^2 = \widehat{\sigma}_{c,t-1}^2 + \gamma\big((e_t^c)^2 - \widehat{\sigma}_{c,t-1}^2\big)\,, \tag{14}$$

where $e_t = x_t - c_t$ is the PLM forecast error.

The agent compares the two models' fit via the time t Gaussian log-likelihood ratio (LLR)

$$\ell_t = \frac{1}{2}\log\frac{\widehat{\sigma}_{t-1}^2}{\widehat{\sigma}_{c,t-1}^2} + \frac{e_t^2}{2\widehat{\sigma}_{t-1}^2} - \frac{(e_t^c)^2}{2\widehat{\sigma}_{c,t-1}^2}. \tag{15}$$

When $\ell_t > 0$, the comparison model fits the current observation better than the PLM.

The complete dynamical system is given by the modified ALM (9), the constant gain recursion for the PLM (8), the constant gain updating of the comparison model and forecast error variances (10)-(14), and the vigilance accumulation (7). Relative to the standard adaptive learning framework, the novel objects are (7) and (15).

3.6 Real time learning and vigilance dynamics

Now we present a set of typical real time learning simulations, comparing the dynamics under steady-state learning to those of vigilance learning. Parameterize the model as $\alpha_L =$

$-1, \alpha_H = 2, \alpha_M = 0.5, k = 5, \sigma = 0.65, \kappa = 20$. The value of κ is based on the finite-gain frontier in Section 5.

Figure 2 plots the results from a long simulation. The model is initialized near the E-stable α_H with a constant gain $\gamma = 0.025$. The left panel (2a) presents a typical least-squares learning simulation in a model with multiple steady state equilibria and a small gain. The learning dynamics for c_t are clustered tightly around the outer stable RPE.[15]

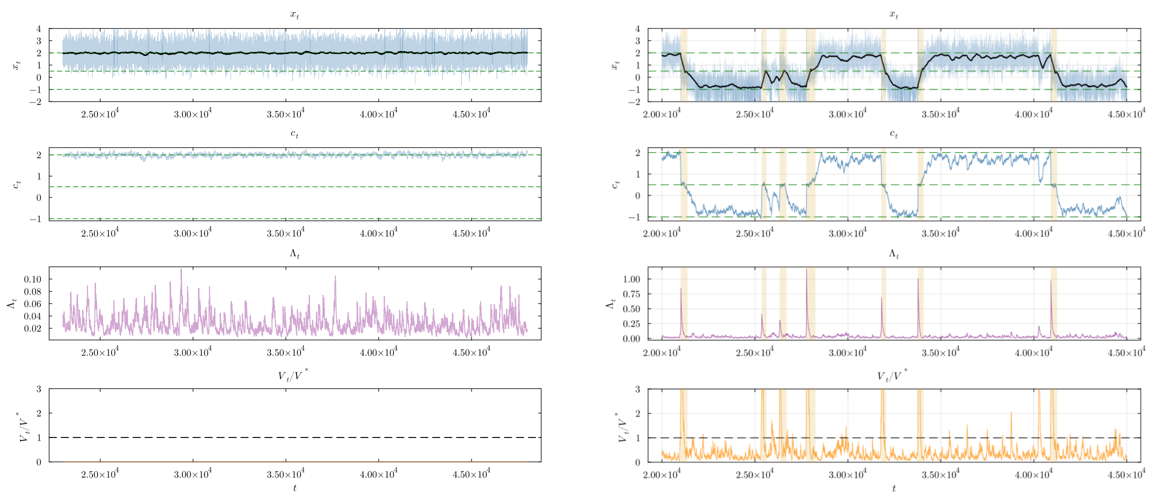

(a) Learning ($\kappa = 0$)

(b) Vigilance learning ($\kappa = 20$)

Figure 2: Real time learning dynamics. The solid black line is a 250 period moving average of x_t.

The right panel (2b) plots the results from a long simulation when the learning has been augmented with vigilance and an endogenous Jensen adjustment. Learning, in this case, displays qualitatively different dynamics. One difference is that the learning dynamics are not contained to the α_H outer well. Learning occasionally crosses into the basin of attraction of the other outer well.

However, the more interesting and distinct behavior occur during those times when (1.) c_t enters and briefly lingers at the unstable α_M equilibrium, which (2.) coincides with spikes in vigilance Λ_t. The periods of temporary lingering at the unstable equilibrium is apparent also in the top panel, where the moving average of x_t repeatedly leaves a stable outer well and approaches the middle steady state α_M.

[15]This behavior is consistent with Evans and Honkapohja (2001) who derive the asymptotic distribution for constant gain learning and show that it converges weakly to a Gaussian whose mean is the E-stable equilibrium with a variance proportional to the gain γ.

An event study analysis makes this novel and puzzling dynamics even more apparent. Figure 3 divides that long simulation into events ($n = 8$) where the learning dynamics begin near the stable outer well α_H and transition to a neighborhood of the unstable α_M and lingers there for at least 50 periods. Figure 3 plots the destination-contingent (α_L, α_H) average across the 8 episodes.

Prior to the departure event ("escape"), c_t and x_t are fluctuating near the upper well. There is, however, a slow accumulation of vigilance in Λ_t as agents doubt their model and apply a small Jensen adjustment. However, that accumulation hits a threshold and there is a sharp rise in vigilance, a sudden Jensen adjustment that reduces the expected value of x_t and, with it, c_t. From there, learning dynamics linger at the *unstable* equilibrium while vigilance remains high. As vigilance relents, the stabilizing attraction of the outer well takes hold and learning dynamics slowly converge back to an outer well.[16]

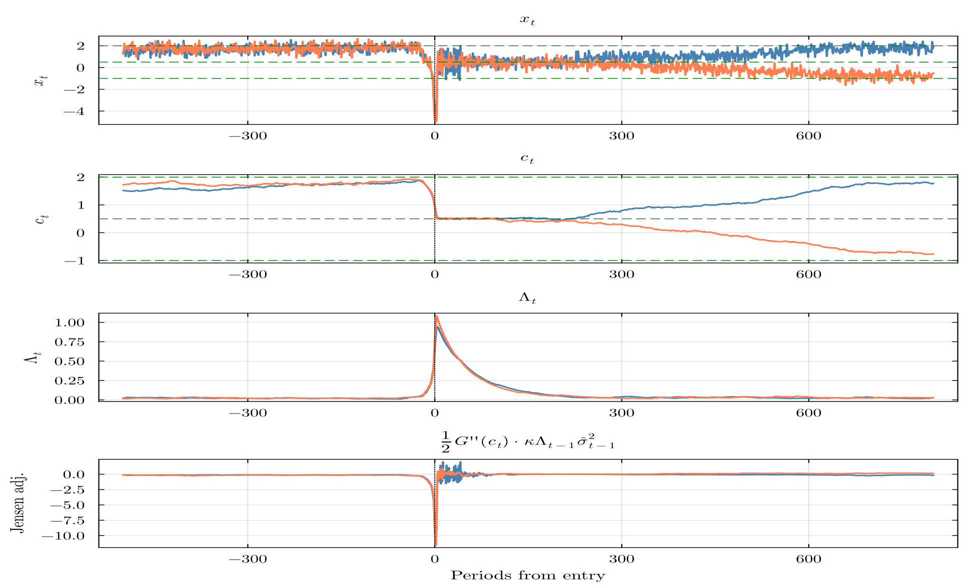

Figure 3: Departing the stable α_H equilibrium and lingering at the unstable α_M equilibrium.

We close this section with a heuristic account of why vigilance learning temporarily stabilizes an unstable equilibrium. On the slow timescale, constant-gain learning separates into two forces. The mean dynamics are a drift that moves down a potential whose two wells lie at the E-stable RPEs, with α_M forming a barrier, while the escape dynamics are rare persistent shock sequences that push beliefs uphill and, in conventional learning, carry

[16]The lingering at the threshold state depends, in part, on the κ and γ, with some values leading to learning-induced noise that keeps vigilance elevated, as demonstrated in Section 5

them over the barrier into the other basin. The key insight of this paper is that vigilance alters the landscape itself. The Jensen adjustment $\frac{1}{2}G''(c_t)V_t$ enters the drift directly, so rising vigilance does not just add volatility. It flattens the wells and, past a critical intensity V^*, collapses them entirely, leaving α_M as the unique attractor. In the cubic-map case the heightened vigilance is naturally constrained. At α_M, $G''(\alpha_M) = 0$ so the Jensen adjustment switches off, α_M is a fixed point for all V, and on either side the Jensen adjustment contracts expected outcomes towards the threshold. At that moment, x_t is i.i.d., the benchmark no longer outpredicts the PLM, vigilance decays, the double well returns, and the temporarily stabilized middle steady state again repels beliefs toward an outer well. Figure 5 at the opening of the next section traces this sequence across the potential landscape.

Figure 4 shows a simulation presenting this cycle. It plots the (centered) belief $\delta_t = c_t - \alpha_M$ against the realized barrier height $\Delta U(V_t)$ along a stretch of the simulation in Figure 2b containing lingering events. The barrier is high in quiescence, collapses to zero exactly when V crosses V^*, stays at zero while beliefs linger at the sticky threshold, and is restored when vigilance recedes. In the middle event beliefs escape the α_L well, linger, and settle at α_H. In the final event they leave α_H, linger, and return to it.

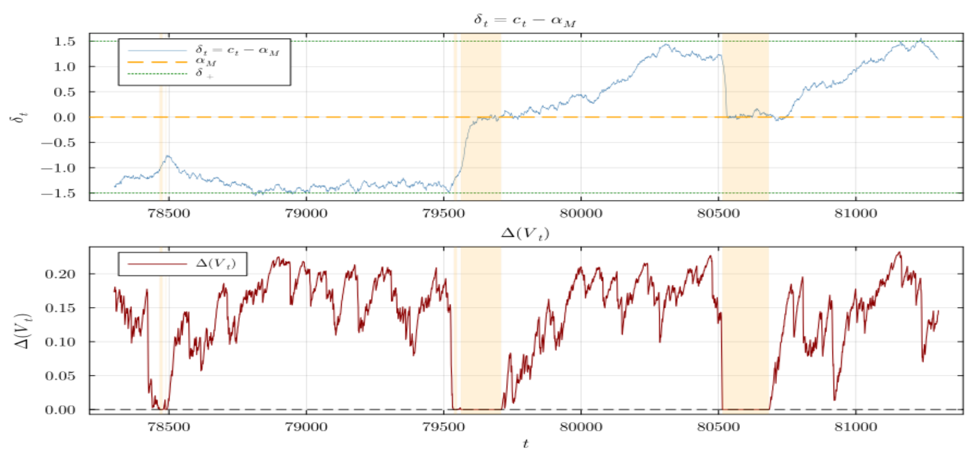

Figure 4: A lingering episode. Top: simulated belief path $\delta_t = c_t - \alpha_M$ during a transition event. Bottom: realized barrier height $\Delta U(V_t)$ along that path. The barrier collapses to zero during the transition, creating the sticky threshold at α_M, and later recovers as vigilance recedes.

The next section formalizes this mechanism, adapting results from the literature on large deviations and small gain approximations.

4 Analytic results

This section establishes one route to a sticky-threshold episode. Ordinary constant-gain learning keeps beliefs near an outer E-stable RPE while vigilance stays small. A rare transition then carries beliefs toward α_M, generating evidence against the perceived law of motion. If the resulting vigilance is high enough for long enough, the outer wells collapse and beliefs linger at α_M. Lingering perists until the benchmark's advantage fades, vigilance recedes, and the double-well returns. Figure 5 depicts the three phases. Release from α_M restores the landscape of the first panel.

The partition makes the formal analysis tractable. Near the RPE, the belief dynamics (c_t) are local. Along the transition only the scalar belief path moves against its mean dynamics, so we condition on a belief profile. Given that path, the benchmark coefficients, second moments, variance estimates, and vigilance-concern stock are, effectively, constant-gain averages of an exogenous signal-plus-noise process. Standard stochastic-approximation arguments show that they track their deterministic mean-dynamics (Evans and Honkapohja, 2001). Large deviations govern the rare belief path while the law of large numbers handles the remaining coefficients. Phases 1 to 3 describe the dynamics conditional on the stage of the episode. The final subsection gives the unconditional probability of a full episode, in the spirit of Cho et al. (2002),Williams (2019), and Williams (2026) by connecting the phases along one complete route with finite exponential cost. The approach does not claim to identify which route is the dominant escape path. The point is to understand the mechanism behind the sticky threshold observed in simulations.

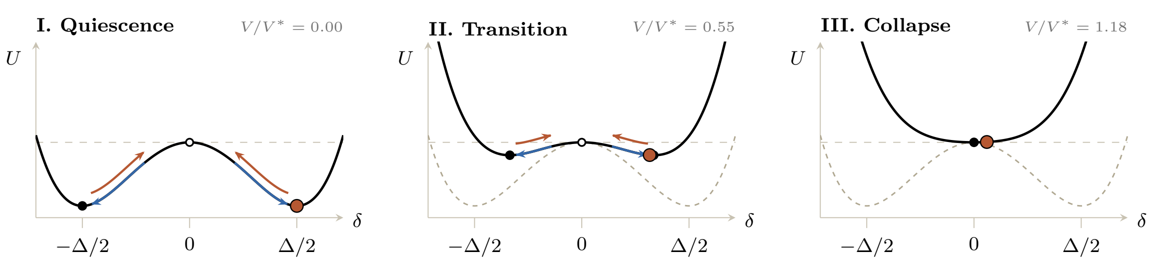

Figure 5: The potential landscape across the phases of a collapse episode. Each panel plots $U(\delta; V) = k^{-1}\left[\delta^4/4 + \frac{3}{2}(V - V^*)\delta^2\right]$ with fixed V ($V = 0$ dashed for reference). Blue arrows denote mean dynamics pressing beliefs down and orange arrows escape dynamics pushing up. $\delta = \pm\Delta/2$ are the $V = 0$ outer wells, filled circles are E-stable steady states, open circle is the E-unstable middle steady state, and orange is $\delta_t = c_t - \alpha_M$.

4.1 Model, notation, and regularity

Define the learning state $Z_t = (c_t, \zeta_t)$, where ζ_t collects the augmented vigilant learning estimates. Since the benchmark regresses on x_{t-1}, the Markov state is $\tilde{Z}_t = (Z_t, x_{t-1})$. It is standard in constant gain learning models to impose some regularity via projection facilities that ensure recursive estimates remain well-behaved by projecting them onto compact sets. We adopt the following assumptions, written here heuristically.

Assumption 1 (Projection)**.** The recursive variance estimates for the PLM and the benchmark model are projected onto $[\underline{\sigma}^2, \overline{\sigma}^2]$, with $0 < \underline{\sigma}^2 < \overline{\sigma}^2 < \infty$. The benchmark's sample second moment matrix is projected onto a compact subset of positive definite matrices. The PLM coefficient c_t, the vigilance state Λ_t, and the benchmark coefficients are projected onto compact sets large enough not to bind on the paths local to the outer wells or on the transition paths considered below. Denote by $\mathcal{Z}$ the resulting compact state space.[17]

Assumption 2 (Lipschitz regularity)**.** On $\mathcal{Z}$, the state-update maps are locally Lipschitz with random Lipschitz constants having uniformly bounded second moments under the Gaussian innovation. Moreover, the LLR (15) is locally Lipschitz in $(c, \vartheta, \widehat{\sigma}^2, \widehat{\sigma}^2_c)$, with a Lipschitz constant uniform over $\mathcal{Z}$:

$$|\ell_t(c, \vartheta, \widehat{\sigma}^2_c, \widehat{\sigma}^2) - \ell_t(c', \vartheta', \widehat{\sigma}^{2\prime}_c, \widehat{\sigma}^{2\prime})| \leq L_t \left(|c - c'| + \|\vartheta - \vartheta'\| + |\widehat{\sigma}^2_c - \widehat{\sigma}^{2\prime}_c| + |\widehat{\sigma}^2 - \widehat{\sigma}^{2\prime}|\right), \quad (16)$$

and $\mathbb{E}[L_t^2] \leq C$, for all parameter values in the projected compact set.

Both conditions are standard. Assumption 1 is the usual projection facility for constant-gain learning. It bounds the state and prevents singularity in the moment matrices. It also supports the local Lipschitz bound (16) in Assumption 2.[18] The standard $O(\gamma)$ mean-square fluctuation of the benchmark estimates around their population values near a stable fixed

[17]Throughout, we use the notational convention that $z \in \mathcal{Z}$, and $z(\tau)$ is the slow-time path satisfying $z(\gamma t) = z_t$. Further, here and the Appendix, asymptotic statements are for small gains $\gamma \to 0$. For a deterministic variable $x_\gamma = O(a_\gamma)$ means that x_γ/a_γ remains bounded, while $o(a_\gamma)$ means that the ratio converges to zero. For random variables, $x_\gamma = O_p(a_\gamma)$ means that x_γ/a_γ is bounded in probability, and $o_p(a_\gamma)$ means the ratio goes to zero in probability. Except when explicitly noted, the constants in these bounds do not depend on γ. So, $O_p(\sqrt{\gamma})$ gives an upper stochastic bound for fluctuations that vanishes as $\gamma \to 0$, while $O_p(1)$ does not vanish.

[18]Note that since x_t depends on the Gaussian $\sigma\eta_t$, the projection facility does not place a bound on x_t. Assumption 1 and differentiating the Gaussian LLR shows that its random Lipschitz coefficient satisfies $L_t \leq C\big(1 + x_t^2 + x_{t-1}^2\big)$. Although (x_t, x_{t-1}) are unbounded, the projection facility bounds their conditional means and the Gaussian innovations provide uniform fourth-moment bounds, implying $\mathbb{E}L_t^2 \leq C$ uniformly over $\mathcal{Z}$.

point, $\mathbb{E}\|\widehat{\vartheta}_t - \vartheta^*\|^2 \leq C\gamma$, is not assumed. It is derived in the proof of Lemma 1(iii). As remarked in the Supplementary Appendix, the Lipschitz bound could equivalently be derived as a lemma but is left here as an assumption to conserve space.

4.2 Pitchfork bifurcation with fixed V

The main intuition for the novel role played by vigilant learning can be seen in the special case where V_t is exogenous and held to a fixed value V, thereby serving as a parameter in the T-map.

To begin, rewrite the constant gain recursion in terms of $\delta_t = c_t - \alpha_M$. For fixed V and small γ, the mean drift and its potential are

$$\mu(\delta; V) = -\frac{\delta[\delta^2 + 3(V - V^*)]}{k}, \quad U(\delta; V) = \frac{1}{k}\left[\frac{\delta^4}{4} + \frac{3}{2}(V - V^*)\delta^2\right],$$
$$V^* = \frac{\Delta^2}{12}, \qquad \mu(\delta; V) = -\partial_\delta U(\delta; V).$$

The derivation is in the proof of Proposition 1. The following proposition provides the mechanism through which vigilance alters the dynamical system's landscape.

Proposition 1 (Fixed-V pitchfork)**.** *Consider the symmetric cubic model with the intercept-only PLM and fixed V. Then,*

1. *For $V < V^*$, the fixed points are $\delta = 0$ and $\delta_\pm(V) = \pm\sqrt{3\,(V^* - V)}$. The middle point is E-unstable, and the two outer equilibria are E-stable.*

2. *For $V > V^*$, $\delta = 0$ is the unique E-stable fixed point.*

3. *For $V < V^*$, the barrier height from either outer well to the middle point is $\Delta U(V) = U(0; V) - U(\delta_+(V); V) = \frac{9(V^* - V)^2}{4k}$.*

With a fixed $V > 0$, $T_V(c)$ flattens as the vigilance parameter increases, and as V crosses the critical value V^* (dependent on $\Delta = \alpha_H - \alpha_L$), the outer wells collapse and the system is left with a unique E-stable RPE at the middle, previously unstable, equilibrium. The far more interesting finding is that endogenous vigilance, with V evolving in real time, leads to episodes where the outer wells are temporarily destroyed and the system lingers near α_M as a sticky threshold. That is the focus of the remainder of this section.

4.3 Phase 1: quiescence

Quiescence is the stationary regime local to an outer E-stable RPE. Standard constant-gain learning keeps beliefs within an $O_p(\sqrt{\gamma})$ neighborhood of that RPE, i.e. small fluctuations that vanish as $\gamma \to 0$. In the fixed-V local linear approximation, their slow movement produces weak serial correlation in the data, that is, $c_t - c^* = O_p(\sqrt{\gamma})$, and $\rho_\gamma \equiv \mathrm{Corr}(x_t, x_{t-1}) = O(\gamma)$. The intercept-only PLM is, near an E-stable RPE, nearly correct, but its best AR(1) comparison detects an $O(\gamma)$ autocorrelation. At the population-optimal benchmark coefficients, the Gaussian LLR gives $\mathbb{E}\ell_t^{\text{pop}} = O(\gamma^2)$, and $\mathbb{E}(\ell_t^{\text{pop}})^+ = O(\gamma)$. This is a key result. While the average population LLR advantage to the benchmark is second-order, the truncated LLR, where concern only accumulates with evidence against the PLM, is first-order.

With recursively estimated benchmark coefficients, on the other hand, Assumptions 1 and 2 produce the weaker bound $\mathbb{E}(\ell_t^{\text{est}})^+ \le C\sqrt{\gamma}$, which, in turn, is shared by Λ_t, hence V_t.

Proposition 2 (Quiescent vigilance and V scaling)**.** *Under Assumptions 1 and 2, for fixed κ and in the local stationary regime near an outer E-stable RPE, as $\gamma \to 0$,*

$$\Lambda_t = O_p(\sqrt{\gamma}), \qquad V_t = \kappa \Lambda_{t-1} \widehat{\sigma}^2_{t-1} = O_p(\sqrt{\gamma}), \qquad \Pr(V_t < V^*) \to 1.$$

Thus, at any date within Phase 1, vigilant learning keeps misspecification concern small and preserves the original double-well potential with probability approaching one. Phase 1 ends only through a rare escape, the subject of Sections 4.4 and 4.6.

4.4 Phase 2: transition

Phase 2 asks what changes when beliefs depart from the quiescent neighborhood, by conditioning on a gradual departure and the benchmark and vigilance respond through their constant-gain averages. As c_t moves toward α_M, the intercept-only PLM cannot track the moving conditional mean in x_t, whose drift creates comovement that the AR(1) benchmark can exploit. The accumulated expected positive LLR is then bounded away from zero. If κ is sufficiently large, V crosses V^* and destroys the outer wells via the pitchfork bifurcation rather than carrying beliefs across the separatrix.

Defining a transition. Define the slow time $\tau = \gamma t$ and define an escape path from an outer RPE by an absolutely continuous profile ψ. For most paths, the slow-time interpolation follows the mean dynamics up to $O_p(\sqrt{\gamma})$ fluctuations. A transition, though, arises from a persistent sequence of unusual shocks that forces beliefs to move counter to their locally stabilizing dynamics. The c_t recursion, rewritten, is $x_t = c_t + \frac{c_{t+1}-c_t}{\gamma}$, so x_t depends on the level and first-difference of beliefs. As $\gamma \to 0$ the deterministic component of the observed signal in slow-time is $\xi_\psi(\tau) = \psi(\tau) + \dot{\psi}(\tau)$. The benchmark observes a noisy signal whose conditional mean moves with ψ_ψ. This distinction is important because model comparison comes from the observed data x_t, and its deterministic component derives from both the level and the speed of the belief transition. The AR(1) benchmark does not need to align with the true law of motion. It just needs to detect this moving mean.

At the transition's outset, the intercept drift $T_V(c) - c$ equals $b_0(c) = G(c) - c$ up to an $O_p(\sqrt{\gamma})$ Jensen adjustment as $\gamma \to 0$ (Prop. 2). A gradual escape toward α_M must climb the wall of the outer well against that stabilizing force. In the scalar case, the least-cost way to escape that well is to cancel the drift via a reverse drift with the opposite sign, that is,

$$\dot{\psi}(\tau) = -b_0(\psi(\tau)) = \psi(\tau) - G(\psi(\tau)). \tag{17}$$

With $V = 0$, this is the least-cost uphill path (Williams, 2026, Proposition 3), the classical time reversal of the deterministic drift (Freidlin and Wentzell, 1998, Ch. 4.3). Along it, $\xi_\psi(\tau) = 2\psi(\tau) - G(\psi(\tau)) \equiv H(\psi(\tau)),\ H(c) = 2c - G(c)$. The reverse-drift restriction makes the econometric signal, observed by the benchmark, a fixed transformation of beliefs.

Definition 3 (Reverse-drift transition path)**.** Fix a sufficiently small endpoint truncation $\varepsilon > 0$. A path $\psi : [0, \Delta\tau] \to \mathbb{R}$ belongs to the *admissible class* $\mathcal{A}$ if it solves (17) and $(\psi(0), \psi(\Delta\tau)) \in \{(\alpha_H - \varepsilon, \alpha_M + \varepsilon), (\alpha_L + \varepsilon, \alpha_M - \varepsilon)\}$.

The path begins once beliefs leave a small neighborhood of the outer RPE and ends when they enter a small neighborhood of the threshold.[19] On either interval between α_M and an outer RPE, $G(c) - c$ has a fixed sign, so ψ moves monotonically toward α_M. Critically, though, the observed signal used to detect the transition, $H \circ \psi$, need not be monotone.

[19]Truncating the endpoints makes $\Delta\tau$ finite, otherwise the outer RPE and α_M are fixed points of the reverse-drift ODE and a path between them cannot occur in finite slow time.

Benchmark detectability. We now address whether the transition generates a benchmark advantage, after a burn-in period for the benchmark coefficients. Conditional on ψ, the transition leads to the signal $X_j^{\psi} = \xi_{\psi}(j\gamma) + \sigma\eta_j, \eta_j \sim N(0,1)$, which, unlike x_t, is exogenous. The approach here is tractable because we do not need to identify any other rare path.

Three uniform properties of $\xi_{\psi} = H \circ \psi$ are key (Lemma 2(i) in the Appendix). First, the signal is L_{ξ}-Lipschitz, so its path cannot be put into an arbitrarily short time interval. Second, its displacement between endpoints is at least $D_{\xi}^{\min} = |\alpha_o - \alpha_M| - 2L_H\varepsilon > 0$, where $\alpha_o \in \{\alpha_L, \alpha_H\}$, implying non-trivial movement in the mean. Third, it has no more than three monotonicity intervals because G' is quadratic. The signal may reverse direction, but not often enough to obscure its movement from the benchmark.

The benchmark coefficients initialized following Phase 1 require a slow-time burn-in τ_a and a residual-window condition: $\tau_a \geq \tau_a^*,\ D_{\xi}^{\min} > L_{\xi}\tau_a$. If that condition holds, the transient period is long enough but there is still a positive amount of displacement remaining. The first inequality ensures the benchmark adapts beyond its quiescent initial condition, while the second prevents the burn-in period from accounting for all of the signal movement. With finite monotone stretches some residual displacement fits into a fixed-length window J_{ψ}, occurring after burn-in. On a sub-block $K \subset J_{\psi}$, signal movement generates a strictly positive population AR(1) advantage. The Appendix then lifts that local population advantage to the mean-dynamics. Let $\bar{u}(s;\psi)$ denote the mean-dynamics expected one-sided LLR under the transition's law of motion: $\bar{u}(s;\psi) \equiv \mathbb{E}_{\eta}\big[\max\{\ell^{\psi}(s), 0\}\big]$. The lemma proves that for constants $h, u_0 > 0$ independent of ψ and γ, $\int_{J_{\psi}} \bar{u}(s;\psi)\, ds \geq hu_0$, uniformly over $\psi \in \mathcal{A}$. Vigilance exponentially weights this accumulation, and Proposition 5 in the Appendix establishes that

$$\ell_*(\psi) \equiv \int_0^{\Delta\tau} e^{-(\Delta\tau - s)}\bar{u}(s;\psi)\, ds \geq e^{-\Delta\tau} hu_0 \equiv \underline{\ell} > 0$$

uniformly over $\psi \in \mathcal{A}$. The key conclusion from this is that the accumulated evidence for the benchmark is bounded away from zero even as $\gamma \to 0$.

Crossing the threshold. The remaining step shows that, conditional on ψ, the benchmark and vigilance recursions track the accumulation just described. The following theorem presents the key result on misspecification concern accumulation during Phase 2.

Theorem 1 (Transition raises concern)**.** *Fix $\psi \in \mathcal{A}$ and, conditional on that profile, let $X_j^{\psi} = \xi_{\psi}(j\gamma) + \sigma\eta_j$ over $[t_a, t_b]$, where $t_b = t_a + N$ and $N = \lfloor \Delta\tau/\gamma \rfloor$. Under Assumptions 1 and 2, suppose the endpoint truncation is small enough for Lemma 2(i), the residual-window condition holds, and the tracking conditions of Lemma 2(iii) hold. Then*

$$\Lambda_{t_b} = e^{-\Delta\tau}\Lambda_{t_a} + \ell_*(\psi) + o_p(1), \qquad \ell_*(\psi) \geq \underline{\ell} > 0, \tag{18}$$

where both the lower bound and the $o_p(1)$ remainder are uniform over $\psi \in \mathcal{A}$. If the transition begins from the Phase 1 local stationary regime, then $\Lambda_{t_a} = O_p(\sqrt{\gamma})$ and hence $\Lambda_{t_b} = \ell_*(\psi) + o_p(1),\ \Pr(\Lambda_{t_b} \geq \underline{\ell}/2) \to 1$.

Phase 1 implies that at the onset of a transition event, concern is small, mainly reflecting the noise floor associated to learning. That low initial concern decays because of constant gain updating. But Phase 2 creates comovement that is detectable by the benchmark. Theorem 1 shows that there is a lower bound $\underline{\ell}$ on the deterministic accumulated LLR. The following corollary converts this into a sufficient condition on κ for crossing $V > V^*$.

Corollary 1 (Sufficient pitchfork-crossing condition)**.** *Suppose the transition begins from the Phase 1 local stationary regime and satisfies the hypotheses of Theorem 1. If $\widehat{\sigma}_t^2 \geq \underline{\sigma}^2 > 0$ and $\kappa > \kappa_C \equiv 2V^*/(\underline{\sigma}^2\underline{\ell})$, then, uniformly over $\psi \in \mathcal{A}$, $\Pr(V_{t_b+1} > V^*) \to 1$.*

The result is immediate from Theorem 1 and Assumption 1 since $V_{t_b+1} = \kappa\widehat{\sigma}_{t_b}^2\Lambda_{t_b} \geq \kappa\underline{\sigma}^2\underline{\ell}/2 > V^*$. The critical value κ_C shows that a larger variance floor $\underline{\sigma}^2$, a more detectable transition (larger $\underline{\ell}$), or a lower pitchfork threshold V^* requires less sensitivity κ to collapse the wells.[20]

4.5 Phase 3: collapse and lingering

Phase 2 shows that a transition can raise the Jensen intensity above V^*. Phase 3 describes events where vigilance is high and persistent enough that the pitchfork lasts long enough to bring c near α_M. Once that happens the separatrix becomes a temporary attractor—a sticky threshold. At that temporary attractor the system lingers (Phase 3a), the intercept-PLM and benchmark AR(1) become indistinguishable generating a decline in vigilance (Phase 3b), the outer wells restore, and the system escapes from the sticky threshold (Phase 3c).

[20] The cubic's advantage is that it bounds the number of monotone stretches of the observed signal to three (uniformly over the admissible paths). The argument works for any smooth G where the number of monotone stretches of $H \circ \psi$ is bounded (uniformly over $\mathcal{A}$) by a finite $M(G)$, after an adjustment of the residual condition.

Definition 4 (Initiating states)**.** Fix $0 < \nu < V^*$ and the fixed small neighborhood of α_M specified in the Appendix. The set of *initiating states* $\mathcal{I}_C^\nu$ consists of states z at which (i) $V(z) \geq V^* + \nu$ and (ii) the deterministic mean dynamics from z carry beliefs into that neighborhood before the Jensen intensity falls below $V^* + \nu$.

Three facts support this dual role of α_M. *(F1) Fixed center.* Since $G(\alpha_M) = \alpha_M$ and $G''(\alpha_M) = 0$, $T_V(\alpha_M) - \alpha_M = 0$ for every V. *(F2) Stability reversal.* Locally, the drift rate is $\lambda(V) = \frac{3}{k}(V^* - V)$, with α_M locally stable when $V > V^*$ and repelling when $V < V^*$. Outside the pitchfork neighborhood $[V^* - \nu, V^* + \nu]$, $|\lambda(V)| \geq \lambda^* \equiv 3\nu/k$: $-\lambda(V)$ is the contraction rate above, and $\lambda(V)$ the repulsion rate below, the neighborhood. *(F3) Losing detectability.* Since $H(\alpha_M) = \alpha_M$, the state equals α_M plus i.i.d. shocks. In a corresponding RPE, the AR(1) benchmark and intercept-only PLM have identical predictive densities and the population LLR benchmark vanishes pointwise. So, while the Phase 2 transition generated drift toward the sticky threshold, detected by the benchmark, that comovement vanishes after arrival into Phase 3.

Let τ_0 be the slow time of entry into $\mathcal{I}_C^\nu$. Fix the escape radius $\varepsilon_{\text{esc}} > 0$ as in the Appendix. Now define the following slow-times

$$
\begin{aligned}
\tau_+ &:= \inf\{\tau \geq \tau_0 : V(\tau) \leq V^* + \nu\}, \\
\tau_- &:= \inf\{\tau \geq \tau_0 : V(\tau) \leq V^* - \nu\}, \\
\tau_{\text{esc}} &:= \inf\{\tau \geq \tau_- : |c(\tau) - \alpha_M| \geq \varepsilon_{\text{esc}}\},
\end{aligned}
\tag{19}
$$

with $\inf \emptyset = +\infty$. τ_+ is entry into the pitchfork neighborhood, while τ_- is its exit, or release, time. τ_{esc} is the escape time from the sticky threshold. The sticky-threshold lingering time is, thus, $\tau_{\text{linger}} = \tau_{\text{esc}} - \tau_0$.

3a. Pinning. The longer V stays above $V^* + \nu$, the greater the contraction toward α_M. So, the key state variable for pinning the system at the sticky threshold is the episode's vigilance budget:

$$
A^-(\gamma) \equiv \int_{\tau_0}^{\tau_+} \frac{3}{k}\big(V(s) - V^*\big)\, ds. \tag{20}
$$

Every marginal unit of budget allows c to continue contracting toward α_M. By the time V reaches the pitchfork neighborhood (at τ_+), the beliefs have a cumulative linear contraction toward α_M by at least a factor $e^{-A^-(\gamma)}$. Constant-gain shocks continue at the $\sqrt{\gamma}$ scale needed for the Phase 3c escape. The Appendix shows that the release (from pitchfork) scale at τ_- is the larger of $e^{-A^-(\gamma)}$ or $\sqrt{\gamma}$, i.e. $O_p(e^{-A^-(\gamma)} \vee \sqrt{\gamma})$. Crossing through the pitchfork neighborhood to release (at τ_-) preserves this bound.

3b. Vigilance decline. With pinning, the population signal that generated misspecification concern vanishes. The benchmark estimates coming from Phase 2 revert over bounded slow time. Then, from its initial value in this phase, Λ declines exponentially with the LLR terms remaining small. Consequently, V crosses below $V^* - \nu$ within a random crossing time that is $O_p(1)$ slow-time. This decline impacts how long collapse persists and eventually the system exits the pitchfork neighborhood.

3c. Escape. Once V first crosses below $V^* - \nu$, instability at α_M is restored and begins to repel the residual distance from the threshold along with new constant-gain shocks. The escape need not be monotone. Movement away from α_M can briefly rebuild vigilance, and vigilance can temporarily flatten the map or even re-contract beliefs. What matters for eventual exit from the ε_{esc}-neighborhood is the cumulative amount of this renewed drag. Define the post-release *exit-drag budget* by $A^+_{\text{exit}}(r) \equiv \int_{\tau_-}^{\tau_-+r} \frac{3}{k} V(s)\, ds$ for $r \geq 0$. The total linear repulsion satisfies $\Phi^*(r) \equiv \int_{\tau_-}^{\tau_-+r} \lambda(V(s))\, ds = \frac{3V^*}{k} r - A^+_{\text{exit}}(r)$. as renewed concern draws from the exit-drag budget. The Appendix shows that this budget is bounded in probability before beliefs leave the ε_{esc}-neighborhood of α_M. Escape begins gradually, giving little sign of drift for most of the window, and then too large and briefly near the end, to replenish vigilance indefinitely. As a result, renewed concern can delay the exit by only a bounded amount in the repulsion exponent.

If the release scale, the larger of $e^{-A^-(\gamma)}$ and the noise floor $\sqrt{\gamma}$, is d, amplification to a fixed radius takes approximately $(k/(3V^*)) \log(1/d)$ units of slow time, up to the bounded exit drag. The earlier stages remain bounded in slow time, while escape can take longer as the gain falls, but at most logarithmically. Renewed concern after release only delays the exit by the bounded exit-drag budget. A larger vigilance budget allows for more lingering, but constant-gain learning noice supplies new $\sqrt{\gamma}$ displacements. The now restored replusion propagates this noise to the escape radius, with probability approaching one. With a fixed initial state, the entry budget $A^-(\gamma)$ is bounded and the realized duration is governed by that budget through (21). Lingering near α_M occurs after its original cause begins to fade. After release, renewed concern can delay exit but has a finite drag A^+_{exit}. This gives the episode a budget-in (depth of pinning), budget-out (delay escape) structure. Theorem 2 combines the three stages.

Theorem 2 (Collapse-and-linger episode)**.** *Under Assumptions 1 and 2, fix an initiating state $z_0 \in \mathcal{I}_C^{\nu}$ and start the post-initiation process from $z(\tau_0) = z_0$, with z_0 fixed as $\gamma \to 0$. The conclusions are uniform over fixed compact subsets on which Definition 4(ii) holds with uniform slack. Then $\tau_- - \tau_0 = O_p(1)$, $\tau_{\text{esc}} < \infty$ with probability approaching one, and the post-release exit-drag budget satisfies $A^+_{\text{exit}}(\tau_{\text{esc}} - \tau_-) = O_p(1)$. For $\tau_{\text{linger}} \equiv \tau_{\text{esc}} - \tau_0$,*

$$\frac{k}{3V^*} \min\Big\{A^-(\gamma), \tfrac{1}{2}\log(1/\gamma)\Big\} - O_p(1) \le \tau_{\text{linger}} \le \frac{k}{6V^*}\log(1/\gamma) + O_p(1). \tag{21}$$

In real time, the episode lasts $\tau_{\text{linger}}/\gamma$ periods, so a bounded lingering in slow time can correspont to a long duration when the learning gain is small.

Corollary 2 (Sharp lingering on the noise-floor branch)**.** *Assume the hypotheses in Theorem 2 and fix $C_A \ge 0$. If $A^-(\gamma) \ge \frac{1}{2}\log(1/\gamma) + C_A$, then the vigilance budget is sufficient to pin c to the learning-noise floor. As $\gamma \to 0$, the bounds in (21) collapse to $\tau_{\text{linger}} = \frac{k}{6V^*}\log(1/\gamma) + O_p(1)$, in the eventwise characterization provided below.*

With the fixed projection bounds, $A^-(\gamma) = O_p(1)$ while $\frac{1}{2}\log(1/\gamma) \to \infty$, so with fixed k and small gains, the noise-floor branch is rare. The corollary is eventwise and establishes that whenever an episode accumulates enough vigilance to pin beliefs to the noise-floor, then there is a sharp expression for τ_{linger}. At finite gains, though, such episodes can remain economically relevant. The general theorem allows for the case of an episode to enter $\mathcal{I}_C^{\nu}$ arbitrarily close to α_M that is not a consequence of the entry budget. A separated-entry condition rules out that head start and sharpens the result.[21]

Corollary 3 (Budget-limited episodes under separated entry)**.** *Under the hypotheses of Theorem 2, suppose also that $|c(\tau_0) - \alpha_M| \ge d_{\text{sep}}$ for a constant $d_{\text{sep}} > 0$ independent of γ. Then $\tau_{\text{linger}} = \frac{k}{3V^*}\min\Big\{A^-(\gamma), \frac{1}{2}\log(1/\gamma)\Big\} + O_p(1)$.*

Phase 3 describes a sticky threshold. The phase begins with a vigilance stock that pins beliefs near the formerly unstable α_M. As the PLM and benchmark become similar, concern declines and the outer wells reemerge. The restored repelling force then unwinds each log unit of that depth in $\approx k/(3V^*)$ units of slow time.[22] The system then exits the episode

[21] Separated entry is natural for an episode which arrives from an outer well. It is an additional entry condition and is not implied by Definition 4.

[22] Under fixed projection bounds and bounded release time, $A^-(\gamma) = O_p(1)$ while $\frac{1}{2}\log(1/\gamma) \to \infty$, so the eventually it is the budget branch that binds in the fixed-κ/small-γ limit. With a finite gain either branch may apply.

and the threshold resumes as a separatrix.

4.6 Can collapse occur?

Phases 2–3 characterize collapse and lingering conditional on entry into the collapse channel. This section addresses the question of existence by establishing a finite-cost reachability result. We impose a sample-path large-deviation structure motivated by Williams (2026, Theorem 2).[23] The assumption is used only to interpret the cost of an explicit candidate path as a rate-function value.

Assumption 3 (Sample-path large deviations)**.** The slow-time interpolation of the projected augmented process satisfies a sample-path large-deviation principle on compact horizons with rate function $I_T(z) = \frac{1}{2}\int_0^T q(\tau)^2\, d\tau$, where $q(\tau)$ is the deterministic perturbation of the Gaussian innovation ("unlikely sequence of shocks") that implements path z.

Proposition 3 (A finite-cost complete sticky-threshold episode)**.** *Fix $\rho > 0$ such that the speed-ρ signal $H_\rho(c) = (1+\rho)c - \rho G(c)$ satisfies the analogue of Phase 2's residual-window condition, and let $\underline{\ell}_\rho > 0$ be the resulting accumulated-LLR lower bound. If $\kappa > \kappa_C^\nu \equiv \frac{V^*+\nu}{\underline{\sigma}^2\underline{\ell}_\rho}$, then there are compact neighborhoods K_o of an outer RPE and $K_C \Subset \mathcal{I}_C^\nu$, and a complete-episode event $\mathcal{E}_\gamma$ whose probability is at least $e^{-\left(I_C^{\mathrm{cand}}+o(1)\right)/\gamma}$ uniformly over initial states in K_o, where $I_C^{\mathrm{cand}} < \infty$ is the action of the candidate path. On $\mathcal{E}_\gamma$ the process reaches K_C, contracts toward the sticky threshold, vigilance declines endogenously, and beliefs exit the $\varepsilon_{\mathrm{esc}}$-neighborhood of α_M. The probability bound does not require Assumption 3. Given that assumption, I_C^{cand} is an admissible rate-function value.*

The candidate path is the speed-ρ version of the reverse drift $\dot{\psi} = \rho\{\psi - G(\psi)\}$ toward α_M.[24] The unusual shocks that move beliefs uphill also generate the drifting signal of Phase 2. The AR(1) benchmark detects that comovement, concern accumulates to at least $\underline{\ell}_\rho$, and the variance floor yields $V \geq \kappa\underline{\sigma}^2\underline{\ell}_\rho > V^* + \nu$. The endpoint can be chosen in a compact interior subset K_C of the initiating set with uniform slack.

[23] A related result is Williams (2019, Theorem 4.1), whose full conclusions require bounded shocks. Extending those previous regularity conditions to the augmented vigilance recursion is beyond the paper's scope.

[24] At $\rho = 1$ the candidate is the usual scalar reverse drift.

5 Robustness and Applications

The previous analysis was developed in the special case of a symmetric temporary equilibrium map, i.i.d. shocks, and steady-state learning. This section explores robustness to departures from symmetry, i.i.d. shocks, and small gains via two further forms of threshold stickiness. Persistent misspecification, say from serially correlated shocks, can produce non-trivial vigilance in equilibrium. With finite gains, there can be sufficient learning volatility to anchor beliefs at the threshold for long stretches of time. For economic interpretation, these departures are developed in applications to growth cycles and dynamic coordination games.

5.1 Symmetry breaking and finite-gain vigilance

Symmetry. Stability reversal for fixed V does not hinge on symmetry. Define the vigilance-adjusted drift term as $F(c,V) = G(c) + \frac{1}{2}G''(c)V - c$ and assume $G \in C^3$ has an interior fixed point c_M with $G(c_M) = c_M$, $G''(c_M) = 0$, $G'(c_M) > 1$, and $G'''(c_M) < 0$. Then c_M is a fixed point for every V and local stability reverses at $V_c = \frac{2[G'(c_M)-1]}{-G'''(c_M)}$ The dynamics at α_M repel from below V_c and attract above it.[25]. The vigilance mechanism, thus, only needs an unstable c_M occurring at an inflection with $G'''(c_M) < 0$ and sufficiently high vigilance. When G is symmetric, the inflection and threshold points coincide. Without symmetry, high vigilance leads to an imperfect pitchfork bifurcation (cf. Golubitsky and Schaeffer, 1979).

Proposition 4 (Asymmetric unfolding)**.** *For the three-root cubic without the symmetry restriction, let $y = c - \bar{\alpha}$ with $\bar{\alpha} = (\alpha_L + \alpha_M + \alpha_H)/3$ and $(c-\alpha_L)(c-\alpha_M)(c-\alpha_H) = y^3 + p_0 y + q_0$. The fixed-$V$ drift is $F(c,V) = -\frac{1}{k}\{y^3 + (p_0 + 3V)y + q_0\}$. If $q_0 \neq 0$, a fold bifurcation occurs at $V_{\text{fold}} = -\frac{p_0}{3} - \left(\frac{q_0^2}{4}\right)^{1/3}$. For $V < V_{\text{fold}}$ the two outer roots are locally stable and the middle root is unstable. At $V = V_{\text{fold}}$ the middle root and one outer root merge. When $V > V_{\text{fold}}$ the unique real root is locally stable.*

The case $q_0 = 0$ was covered in Proposition 1. With symmetry, the threshold exists with $G(\alpha_M) = \alpha_M, G''(\alpha_M) = 0$ for every fixed V (see Figure 6), playing the role of a separatrix for low V and an attractor for large V. Without symmetry, as V increases,

[25] Because $G''(c_M) = 0$, differentiating gives $F_c(c_M, V) = G'(c_M) - 1 + \frac{1}{2}G'''(c_M)V$, which is decreasing in V and crosses zero at V_c.

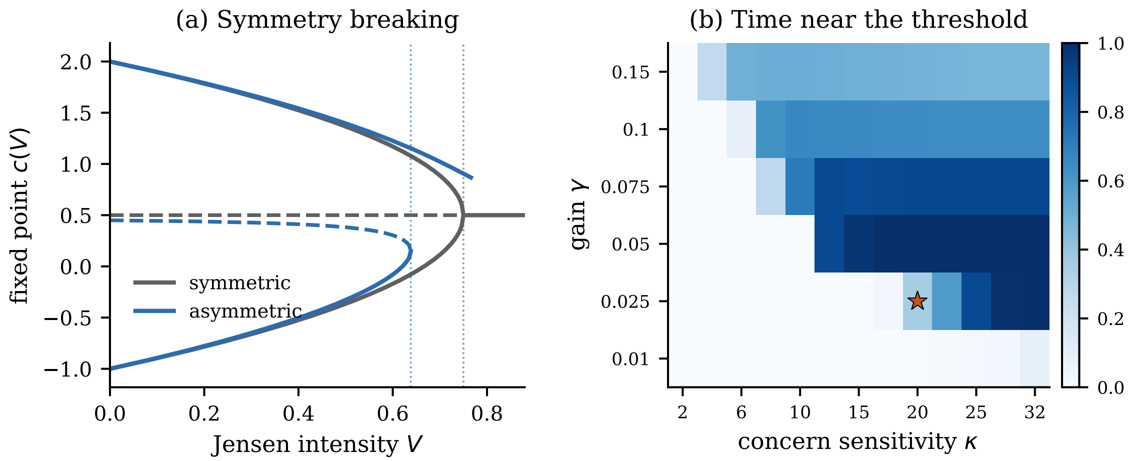

Figure 6: Panel (a) compares the symmetric pitchfork (gray) with a small symmetry-breaking perturbation (blue) (stable branches are solid, dotted lines are critical V). Panel (b) reports the share of a 100,000-period simulation that beliefs spend within 0.35 of α_M at the baseline calibration $\sigma = .65$ (star $(\gamma, \kappa) = (.025, 20)$).

the middle and its nearer outer well fold together at V_{fold}. When $V > V_{\text{fold}}$ there exists a point $y_r(V)$ for $V \in [V_{\text{fold}}, -p_0/3]$ with $3y_r(V)^2 = -(p_0 + 3V)$, a *fold remnant* in the sense of Williams (2026, Definition 3), where the drift's slope is zero and its value is small, proportional to $V - V_{\text{fold}}$.[26] Effectively, the threshold's two roles are now separate. There is still a sticky threshold in the sense that the mean-dynamics are weak near the remnant as the path moves toward the unique stable fixed point.[27] Williams (2026, Corollaries 2–3 and Proposition 7) establishes that escapes concentrate, for this reason, at remnants. The surviving branch in Figure 6 sustains beliefs near α_M as greater V pushes that surviving fixed point toward the center. The system releases, with a bias, once the endogenous V_t recrosses below V_{fold}.[28]

Finite gains. Learning noise combined with a sufficiently high concern sensitivity κ can sustain a high Jensen intensity for extended stretches of time. To see how a fixed $\gamma > 0$ can support high vigilance, hold $c_t = \alpha_M$. Here $G''(\alpha_M) = 0$, so the data are i.i.d., there's no Jensen adjustment, and nothing depends on κ. π^M_γ is the associated stationary distri-

[26] The remnant is not an equilibrium, and it does not appear in panel (a) of Figure 6, which plots the fixed points for $c(V)$.

[27] In a sense, the sticky threshold becomes a sticky region.

[28] The result here is fixed-V only. Verifying that the three phases extend to the asymmetric unfolding remains open.

bution. The following Lyapunov exponent provides a natural diagnostic of local stability: $\chi_M(\gamma,\kappa) \equiv \mathbb{E}_{\pi_\gamma^M} \log \left|1 + \frac{3\gamma}{k}(V^* - V_t)\right|$. A sufficient condition for the fixed-c system to be locally contracting is that $\mathbb{E}_{\pi_\gamma^M} V_t > V^*$.[29] Since $V_t = \kappa \Lambda_{t-1}\widehat{\sigma}_t^2$, the fixed-$c$ system contracts when $\kappa\, \mathbb{E}_{\pi_\gamma^M}[\Lambda\widehat{\sigma}^2] > V^*$. The Phase 1 bounds cap Λ at order $\sqrt{\gamma}$, and so real-time simulations reveal the frontier of the anchored region in panel (b) of Figure 6 tracks $\kappa\sqrt{\gamma}$. Panel (b) of Figure 6 illustrates anchoring in a high vigilance regime by simulating the model for 100,000 periods at $\sigma = 0.65$ over a grid of values for (κ, γ). The figure plots the share of a simulation spent near the sticky threshold. For a fixed gain, the system anchors longer near the threshold for larger values of concern sensitivity κ.[30]

5.2 Growth cycles and persistent population vigilance

Relaxing the i.i.d. shock assumption can lead to an RPE with persistent non-trivial vigilance. Depending on the shock volatility and the concern sensitivity, α_M can continue to play the dual role of a separatrix and attractor, or it can be the conditional mean of a best linear fit to data ranging across both outer wells. We illustrate these features in the endogenous growth model proposed by Evans et al. (1998) (henceforth EHR).

EHR study growth cycles in which adaptive expectations select among balanced-growth paths. Agents forecast returns and in some calibrations there exist E-stable low- and high-growth paths with an E-unstable middle separatrix. We consider their model but choose a technology process so that the temporary equilibrium map fits the cubic form studied in this paper. EHR's temporary equilibrium structure consists of an Euler equation, $g_t = [\beta(1+r_t)]^{1/\sigma}$, and a technology curve relating the real interest rate to expected future growth rates, $r_t = R(g_{t+1}^e)$. We depart from EHR by specifying $R(g^e) = \beta^{-1}\left[G_N(g^e)\right]^\sigma - 1$ in closed form, where G_N is the cubic with roots at EHR's balanced-growth paths and k chosen to match the slope of the temporary equilibrium map at the middle path.[31] Com-

[29]The sufficient condition also requires the term inside the absolute value to be positive almost surely, which holds for small γ on the projected state space. This condition, also, assumes fixed c, i.e. a system that is already near α_M, abstracting from the main feedback and the stability of the joint system. Instead, it delivers a critical value for the concern sensitivity.

[30]However, it should be noted that high values for V in an anchoring regime can lead to large values for x that hit the projection bounds. In simulations it was convenient to truncate $x \in [-5, 5]$.

[31]In EHR the technology curve's slope depends on a balance between complementarity across capital-good types and the rising cost of investment goods. Their formulation has a kinked production possibility frontier, and the implied temporary equilibrium map is locally convex at both outer paths, unlike the cubic used here. The closed form reproduces their equilibria but with the cubic form of the technology curve.

posing these two equations gives the temporary equilibrium $g_t = G_N(g^e_{t+1})$. After a normalization, the balanced-growth paths are set at $(\alpha_L, \alpha_M, \alpha_H) = (-1, .5, 2)$.[32]

Now we depart from i.i.d. shocks and assume that $\varepsilon_t = \rho_\varepsilon \varepsilon_{t-1} + \sigma_u \eta_t$. We no longer have noisy steady-state RPEs and g_t is serially correlated, even in neighborhoods of the E-stable RPEs. Thus, agents use the AR(1) perceived law of Branch and McGough (2005), $g_t = c + b\,(g_{t-1} - c) + e_t$, whose intercept c is the perceived long-run growth rate and whose slope b is the perceived persistence. The benchmark is also extended to detect the misspecification from a linear PLM in a non-linear environment by regressing on a quadratic term that detects curvature away from α_M and a cubic term that spans the S-shape.

Table 1: RPE in EHR at two concern sensitivities. Each row is an RPE that is locally stable under the full mean-dynamics ODE. σ^{2*} and σ_c^{2*} are the population residual variances of the PLM and the benchmark at the RPE.

	κ	c^*	b^*	V_{eq}	V_{eq}/V_c	$\sigma_c^{2*}/\sigma^{2*}$	Λ^*	mean g	sd g
middle (threshold)	700	.503	.875	.759	1.01	.986	.037	.503	.355
	300	.503	.914	.701	.94	.958	.073	.503	.441
outer (high growth)	700	1.957	.305	.042	.06	.9999	.002	1.957	.172
	300	1.983	.300	.017	.02	.9999	.002	1.983	.172

In this richer specification of the model vigilance can remain, in equilibrium, at a nontrivial level that does not vanish with the learning gain, unlike the i.i.d. shock case. Numerically we find evidence for such vigilance RPE. Parameterize the model with $\rho_\varepsilon = .6$ and $\sigma_u = .16$. Table 1 reports the RPE values at high-sensitivity $\kappa = 700$ and the lower $\kappa = 300$.[33] In this growth model with serially correlated shocks three distinct (E-stable) RPE coexist at both values of κ. By symmetry, the outer high-growth RPE (chosen here) has a mirrored low-growth RPE. The expected growth rate is centered at the threshold, as is the actual growth rate with standard deviation $.35$ at $\kappa = 700$ and $.44$ at $\kappa = 300$.[34] Noticing the gap from the outer wells, it's clear that the threshold is a genuine RPE concentrated near there as opposed to a trend line fit to data moving between the outer wells. This

[32]Gross growth is $1.0257 + .0154\,(g_t - .5)$. Numerically it is convenient to apply a centering and unit-rescaling normalization for the regressors. See the Appendix for details on computing the RPE in this application.

[33]While not directly comparable to the baseline i.i.d. case, $\kappa = 300$ is approximately the same as the baseline $\kappa = 20$ after adjusting for the growth model's roughly thirteen-times-smaller forecast variance.

[34]The perceived law is stationary at every RPE, $b^* < 1$, and its implied variance $\sigma^{2*}/(1 - b^{*2})$ reproduces the actual variance, $.0294/(1 - .875^2) = .355^2$ at the $\kappa = 700$ threshold RPE.

sticky threshold co-exists with RPE near the outer wells where misspecification concern is low. Moreover, here κ does not select which RPEs exist. At $\kappa = 700$ the Jensen intensity just crosses the pitchfork critical value, $V_{\text{eq}} > V_c$, and the fixed-V sticky threshold RPE is unique. When $\kappa = 300$, even though $V_{\text{eq}} < V_c$, the sticky threshold equilibrium survives because the PLM and benchmark sample a wider swath of the map's curvature an vigilance concern nearly doubles.[35]

5.3 Smooth strategic coordination and rollover

While the applications so far are macroeconomic, we briefly note that sticky thresholds can also arise in learning in games, as in Fudenberg and Levine (1998, Ch. 4) and Hofbauer and Sandholm (2002). For example, consider a symmetric two-action coordination problem and normalize the payoff difference between action one and action zero to $2p - 1$. The logit response is $B_\beta(p) = \frac{1}{1+\exp[-\beta(2p-1)]}$, with mean-dynamics $\dot{p} = B_\beta(p) - p$. Since $B_\beta(1/2) = 1/2$ and $B'_\beta(1/2) = \beta/2$, the middle state is locally unstable when $\beta > 2$ while there are two stable coordination outcomes. Appendix E shows that the vigilance-adjusted response $B_{\beta,V}(p) = B_\beta(p) + \frac{1}{2}B''_\beta(p)V$ undergoes a pitchfork bifurcation at $V^*_{\text{logit}} = \frac{\beta-2}{\beta^3}$.[36] Below V^*_{logit} the drift pushes participation away from $1/2$, and above it the drift pushes participation toward $1/2$, so vigilance concern stabilizes the formerly unstable threshold.

Interestingly, the coordination environment also has a sovereign-rollover interpretation, e.g. Cole and Kehoe (2000), though of the lender-participation margin rather than their full sovereign-debt model. In this interpretation, p_t is the fraction of lenders rolling over debt, the sovereign's financing needs are subject to random shocks, and there is a strategic complementarity aspect to the decision to rollover debt. Heightened vigilance makes the marginal rollover outcome temporarily stabilizing, acting as a sticky threshold within a crisis region. Thus, the simple framework here can apply in a variety of economic settings including technology adoption, growth cycles, rollover coordination, and learning in coordination games.

[35] Branch and McGough (2005) feature an RPE that arises by fitting a linear trend line to data drawn from a bi-modal distribution. That result can emerge here by fixing V appropriately.

[36] To preserve $p \in [0, 1]$, we also require that $V \leq 4/\beta^2$. The pitchfork is supercritical, as earlier, for $2 < \beta < 8/3$.

6 Conclusion

Many economic environments feature persistent regimes separated by a tipping threshold. This paper asked what happens when the transition reveals that the models agents forecast with are misspecified. Following Lanzani (2025), vigilant learners forecast from a parsimonious model but monitor a richer diagnostic benchmark, and rising concern about misspecification leads them to apply a Jensen adjustment around their point forecast. In a self-referential economy that adjustment feeds back into the law of motion, and when strong enough it reverses the threshold's stability, creating a sticky-threshold. In quiescence near a stable regime, vigilance features small fluctuations. Along a transition, the moving conditional mean is detectable, concern accumulates to a level bounded away from zero, and with enough sensitivity the outer regimes collapse. At the threshold the signal that raised concern disappears, vigilance recedes, and beliefs return to a stable outer well after a lingering time. These episodes occur with a finite-action cost, and they recur under standard recurrence conditions. With persistent misspecification, as in the growth application, the threshold can instead support a population equilibrium with nontrivial vigilance, and at finite gains learning noise alone can anchor beliefs there for long stretches. In each case, what looks like a third regime may be a temporarily stabilized boundary between two. At small gains the sticky threshold is transient while at empirical gains episodes can be frequent or even an equilibrium outcome.

A Algebra and Phase 1 quiescence

This appendix presents the proofs to the main results. In particular, the appendix presents the core lemmas, key mechanisms, and steps to establish the Section 4 results. The Supplementary Appendix provides the proofs of the technical lemmas stated below and the full proof of Proposition 3, together with additional material about the applications.

Proof of Proposition 1. Define $\delta = c - \alpha_M$. Symmetry of the cubic gives $G(c) - c = -\frac{\delta(\delta^2 - \Delta^2/4)}{k}, G''(c) = -\frac{6\delta}{k}$. where $\Delta = \alpha_H - \alpha_L$. Let $V^* = \Delta^2/12$. Then, $\mu(\delta; V) = T_V(c) - c = -\frac{\delta\{\delta^2 + 3(V - V^*)\}}{k}$. For $V < V^*$ the nonzero roots are $\delta_\pm(V) = \pm\sqrt{3(V^* - V)}$, and, $\partial_\delta \mu(0; V) = \frac{3}{k}(V^* - V), \partial_\delta \mu(\delta_\pm; V) = -\frac{6}{k}(V^* - V)$, which gives the stability about the point V^*. Integrating $-\mu$ from zero leads to the potential $U(\delta; V) = \frac{1}{k}\left\{\frac{\delta^4}{4} + \frac{3(V - V^*)}{2}\delta^2\right\}$, $U(0; V) - U(\delta_\pm; V) = \frac{9(V^* - V)^2}{4k}$. $\square$

Lemma 1 (Quiescent local orders)**.** *Fix $V < V^*$ and an outer E-stable RPE $c^* = T_V(c^*)$. Let $\beta = T_V'(c^*)$ and $\lambda = \beta - 1 < 0$. In the local approximation, as $\gamma \to 0$:*

(i) $\mathrm{Var}(c_t) = \frac{\gamma\sigma^2}{2|\lambda|} + O(\gamma^2)$, $\rho_\gamma \equiv \mathrm{Corr}(x_t, x_{t-1}) = \gamma\left(\beta + \frac{\beta^2}{2|\lambda|}\right) + O(\gamma^2)$.

(ii) Suppose that x_t is stationary Gaussian with lag-one autocorrelation ρ. The population LLR compares the best AR(1) and best intercept-only PLMs, and its expected value depends only on (x_t, x_{t-1}). Then the population-optimal AR(1)-versus-intercept LLR satisfies $\mathbb{E}\ell^{\mathrm{pop}} = -\frac{1}{2}\log(1-\rho^2)$, $\mathbb{E}(\ell^{\mathrm{pop}})^+ = \frac{|\rho|}{\pi} + O(\rho^2)$, $\mathbb{E}|\ell^{\mathrm{pop}}| = O(|\rho|)$. In particular, these orders are $O(\gamma^2)$, $O(\gamma)$, and $O(\gamma)$, respectively, when $\rho = \rho_\gamma$.

(iii) Under Assumptions 1 and 2, $\mathbb{E}|\ell_t^{\mathrm{est}}| + \mathbb{E}(\ell_t^{\mathrm{est}})^+ \le C\sqrt{\gamma}$.

Proof of Proposition 2. Taking expectations in (7) under the local stationary distribution gives $\mathbb{E}\Lambda_t = \mathbb{E}\ell_t^+$, and Lemma 1(iii), proved in the Supplementary Appendix for the localized recursion, bounds the right side by $C\sqrt{\gamma}$. Markov's inequality gives $\Lambda_t = O_p(\sqrt{\gamma})$. Assumption 1 then gives

$$0 \le V_t \le \kappa\overline{\sigma}^2\Lambda_{t-1} = O_p(\sqrt{\gamma}), \qquad \Pr(V_t \ge V^*) \le \frac{\kappa\overline{\sigma}^2\mathbb{E}\Lambda_{t-1}}{V^*} \longrightarrow 0.$$

□

B Phase 2: transition detectability

Lemma 2 (Transition mechanics)**.** *Under Assumptions 1 and 2, the following hold uniformly over $\psi \in \mathcal{A}$ for sufficiently small γ and endpoint truncation (Definition 3) satisfying $\varepsilon < \Delta/(4L_H)$.*

(i) The transition signal $\xi_\psi = H \circ \psi$ is L_ξ-Lipschitz, has endpoint displacement at least $D_\xi^{\min} > 0$, and has at most three monotonicity intervals.

(ii) If $\tau_a \ge \tau_a^*$ and $D_\xi^{\min} > L_\xi \tau_a$, there is a window $J_\psi \subset [\tau_a, \Delta\tau]$ of fixed length $h > 0$ and a constant $u_0 > 0$, neither depending on ψ or γ, such that $\int_{J_\psi} \bar{u}(s;\psi)\,ds \ge hu_0$.

(iii) Let $Z_j^{\gamma,\psi}$ be the full-system recursion along $X_j^\psi = \xi_\psi(j\gamma) + \sigma\eta_j$, and let $\bar{Z}^\psi$ be its mean-dynamics path. If $\|Z_0^{\gamma,\psi} - \bar{Z}^\psi(0)\| = O_p(\sqrt{\gamma})$ and the mean path stays in the

nonbinding interior of the projection region, then, for $N = \lfloor \Delta\tau/\gamma \rfloor$, $\sup_{j \le N} \|Z_j^{\gamma,\psi} - \bar{Z}^\psi(j\gamma)\| = O_p(\sqrt{\gamma})$, *and* $\gamma \sum_{j=0}^{N-1} (1-\gamma)^{N-1-j} \{u_j^\psi - \bar{u}(j\gamma;\psi)\} = o_p(1)$.

Proposition 5 (Benchmark advantage)**.** *Under Lemma 2(ii),* $\ell_*(\psi) \equiv \int_0^{\Delta\tau} e^{-(\Delta\tau - s)} \bar{u}(s;\psi)\, ds \ge \underline{\ell} \equiv e^{-\Delta\tau} h u_0 > 0$, *uniformly over* $\psi \in \mathcal{A}$.

Proof of Theorem 1. Iterating (7) over the transition, $\Lambda_{t_b} = (1-\gamma)^N \Lambda_{t_a} + \gamma \sum_{j=0}^{N-1} (1-\gamma)^{N-1-j} u_{t_a+j}$. Start the conditional mean path at the transition's initial state (w.l.o.g. a zero initial gap, else quiescence adds $O_p(\sqrt{\gamma})$). Assumption 1 makes projection nonbinding on this path. Since $(1-\gamma)^N \to e^{-\Delta\tau}$, part (iii) of Lemma 2 and a Riemann-sum approximation give (18). Proposition 5 gives $\ell_*(\psi) \ge \underline{\ell}$. If the transition starts from quiescence, $\Lambda_{t_a} = O_p(\sqrt{\gamma})$, so $\Lambda_{t_b} \ge \underline{\ell}/2$ with probability approaching one. □

C Collapse and lingering

Before presenting the central estimates for this phase a quick note on terminology. We say there is "release" at τ_- whenever the system transits below $V \le V^* - \nu$.

Lemma 3 (Episode estimates)**.** *Condition on entry into the initiating set at* τ_0, *and fix a compact subset of* $\mathcal{I}_C^\nu$ *on which condition (ii) of Definition 4 holds with uniform slack, i.e. the mean-dynamics path for* c *reaches the* $\varepsilon_0/2$ *neighborhood of* α_M, *for any initial state in the subset, within a bounded slow-time while the mean Jensen intensity remains a positive distance above* $V^* + \nu$. *Slack is a property of condition (ii) not* ν. *On such a subset, the continuation construction leads to fixed constants* τ_{flow}, $\overline{T}$, *and* $C_{\text{tube}} > 1$, *independent of* γ. *Fix* $\epsilon_0 > 0$ *so that* $C_{\text{tube}}\epsilon_0$ *is inside the (local) Lipschitz neighborhood and* $\epsilon_0 \le \nu^2$, $\kappa\overline{\sigma}^2 C C_{\text{tube}}\epsilon_0 \le \frac{1}{2}(V^* - \nu)$, *where* C *is the local LLR constant below. There is also a fixed* $\varepsilon_{\text{esc}}^0 > 0$, *independent of* γ, *such that for every fixed* $0 < \varepsilon_{\text{esc}} \le \varepsilon_{\text{esc}}^0$ *the* ε_{esc}*-neighborhood of* α_M *lies inside the nonbinding projection region, the cubic term does not cancel out the linear repulsion there, and movement in* $c-$ *inside it regenerates too little concern after release to keep vigilance from receding. The following estimates hold pointwise in the initial state, and uniformly on any such subset.*

(i) At the upper threshold crossing, $|c(\tau_+) - \alpha_M| = O_p\left(e^{-A^-(\gamma)} \vee \sqrt{\gamma}\right)$.

(ii) There is a fixed $\overline{T}$ *such that* $\Pr(\tau_- - \tau_0 \le \overline{T}) \to 1$, *and the belief remains in the* $C_{\text{tube}}\epsilon_0$ *tube after the initial mean-dynamics interval, lasting* τ_{flow} *in slow time.*

On that tube, define $s \equiv \tau - \tau_0 - \tau_{\text{flow}}$. *Then* $\bar{u}^{\gamma}_M(\tau) \le C(C_{\text{tube}}\epsilon_0 + e^{-s}) + O_p(\sqrt{\gamma})$, $\Lambda(\tau) \le \Lambda(\tau_0 + \tau_{\text{flow}})e^{-s} + Cse^{-s} + CC_{\text{tube}}\epsilon_0 + O_p(\sqrt{\gamma})$, $\bar{u}^{\gamma}_M$ *being the Phase 3 mean-dynamics path for* ℓ^+. *Consequently,* $\tau_- - \tau_+ = O_p(1)$.

(iii) *The pitchfork crossing changes the scale at which* c *is pinned near* α_M *by only a bounded factor. That is, for* $D_\gamma \equiv |c(\tau_-) - \alpha_M| \vee \sqrt{\gamma}$, $D_\gamma = O_p\Big(e^{-A^-(\gamma)} \vee \sqrt{\gamma}\Big)$.

(iv) *For every fixed* $0 < \varepsilon_{\text{esc}} \le \varepsilon^0_{\text{esc}}$, *the post-release exit-drag budget* $A^+_{\text{exit}}(r) \equiv \int_{\tau_-}^{\tau_-+r} \frac{3}{k} V(s)\, ds$ *is bounded in probability through escape, and escape occurs with probability approaching one:* $A^+_{\text{exit}}(\tau_{\text{esc}} - \tau_-) = O_p(1)$, $\Pr(\tau_{\text{esc}} < \infty) \longrightarrow 1$. *Moreover, with* $\Phi^*(r) \equiv \int_{\tau_-}^{\tau_-+r} \lambda(V(s))\, ds = \frac{3V^*}{k} r - A^+_{\text{exit}}(r)$, *we have*

$$\log\frac{\varepsilon_{\text{esc}}}{D_\gamma} - O_p(1) \le \Phi^*(\tau_{\text{esc}} - \tau_-) \le \tfrac{1}{2}\log(1/\gamma) + O_p(1). \tag{22}$$

Consequently, $\frac{k}{3V^*}\left\{\log\frac{\varepsilon_{\text{esc}}}{D_\gamma} - O_p(1)\right\} \le \tau_{\text{esc}} - \tau_- \le \frac{k}{6V^*}\log(1/\gamma) + O_p(1)$. *Recrossings of the pitchfork are not excluded. Their cumulative effect is measured via* A^+_{exit}.

Proof of Theorem 2. Parts (ii)–(iii) of Lemma 3 make the time to release from the pitchfork $[V^* - \nu, V^* + \nu]$ order $O_p(1)$ and also $\log(\varepsilon_{\text{esc}}/D_\gamma) \ge \min\{A^-(\gamma), \frac{1}{2}\log(1/\gamma)\} - O_p(1)$. Part (iv) shows that the episode terminates with beliefs exiting the ε_{esc} neighborhood with probability approaching one and the exit-budget $A^+_{\text{exit}}(\tau_{\text{esc}} - \tau_-) = O_p(1)$. The two inequalities in (21) then follow. Nothing is assumed about vigilance after release and so, for example, a recrossing of the pitchfork is allowed. Adding the bounded time to release changes only the $O_p(1)$ terms. Dividing slow time by γ gives the real-time duration. □

Proof of Corollary 2. Set $L_\gamma \equiv \frac{1}{2}\log(1/\gamma)$ and let $E_\gamma \equiv \{A^-(\gamma) \ge L_\gamma + C_A\}$ be the ample-budget event posited by the corollary. Let H_γ be the event where (21) holds with common non-negative remainders $R_\gamma = O_p(1)$. Theorem 2 establishes $\Pr(H_\gamma) \to 1$. During an event in E_γ the vigilance-accumulation budget is large enough to pin c to the noise floor. Parts (i) and (iii) of Lemma 3 then give a release distance from α_M when $V < V^* - \nu$ that scales like $D_\gamma = O_p(\sqrt{\gamma})$, the minimum in (21) equals its noise-floor branch (L_γ up to C_A), and the lower and upper bounds of Theorem 2 collapse to each other at $\frac{k}{6V^*}\log(1/\gamma)$.

The quantifier behind the collapsing inequality requires care. During an ample-budget E_γ, the min in (21) is $\frac{1}{2}\log(1/\gamma)$. So, on the intersection $H_\gamma \cap E_\gamma$ the theorem's estimates and the ample budget occur simultaneously, so the lower and upper bounds have the same first term, yielding, $H_\gamma \cap E_\gamma \subseteq \left\{ \left|\tau_{\text{linger}} - \frac{k}{6V^*}\log(1/\gamma)\right| \leq R_\gamma + \frac{k}{3V^*}(-C_A)_+ \right\}$. which is the stated result. Under Assumption 1, Lemma 3(ii) makes $A^-(\gamma) = O_p(1)$ while $L_\gamma \to \infty$, so $\Pr(E_\gamma) \to 0$ and conditioning on E_γ selects increasingly unlikely paths. A large vigilance budget requires a long contracting phase because V is bounded on $\mathcal{Z}$. Let $V^{\max} = \max_{z\in\mathcal{Z}} V(z)$. The budget's upper bound is $A^- \leq (3/k)(V^{\max} - V^*)_+(\tau_+ - \tau_0)$. On E_γ, the budget is at least $\frac{1}{2}\log(1/\gamma) + C_A$. So, on E_γ the time before V enters the pitchfork neighborhood at τ_+ diverges as $\gamma \to 0$. The qualification in the corollary is needed since, even though $\Pr(H_\gamma) \to 1$ and $\Pr(E_\gamma) \to 0$, it doesn't necessarily follow that $\Pr(H_\gamma \mid E_\gamma) \to 1$. While H_γ is likely, it doesn't need to occur conditional on the rare event E_γ. So the lingering estimate holds on $H_\gamma \cap E_\gamma$. $\square$

The Supplementary Appendix section "Collapse and lingering" proves Corollary 3 where a lower bound on the product measuring accumulated contracting factors rules out ordinary learning noise cancelling out the budget-limited branch.

D Collapse-channel reachability

Proof of Proposition 3. The full proof is in the Supplementary Appendix section "Collapse-channel reachability." It has four steps, briefly summarized here.

The candidate path has finite cost. Recall $b_0(c)$, let ψ_ρ solve $\dot{\psi}_\rho = -\rho b_0(\psi_\rho)$ from an ε-truncated outer endpoint as in Definition 4(ii). Now though we prepend the path with a smooth connector of fixed duration from the z_+ and remaining inside the quiescent neighborhood, keeping T_ρ finite. Conditional on this c-path, the augmented coefficients follow the mean-dynamics of $\xi_\rho = \psi_\rho + \dot{\psi}_\rho = H_\rho(\psi_\rho)$ plus noise, with implied Jensen intensity $\bar{V}_\rho$. The mean shock $q_\rho = [\xi_\rho - T_{\bar{V}_\rho}(\psi_\rho)]/\sigma$ reproduces $\dot{\psi}_\rho$. It is bounded by the projection facility, so $I_C^{\text{cand}} = I_{\text{conn}} + \frac{1}{2}\int_0^{T_\rho} q_\rho^2\, d\tau < \infty$.

The candidate ends in the interior of $\mathcal{I}_C^\nu$. Following the earlier logic (G' is quadratic, residual-window for H_ρ), the Phase 2 detectability bound (Proposition 5) applies here as $\bar{\Lambda}_\rho(T_\rho) \geq \underline{\ell}_\rho$. With the variance floor and $\kappa > \kappa_C^\nu$, $\bar{V}_\rho(T_\rho) \geq \kappa\underline{\sigma}^2\underline{\ell}_\rho > V^* + \nu$, while $\psi_\rho(T_\rho) \in B_M$. Continuity of the mean dynamics in the initial state gives a compact ball $K_C \Subset \mathcal{I}_C^\nu$ around the terminal state with the entry slack required by Theorem 2.

We then show that the constructed tube around the candidate is exponentially likely.[37] Let P_γ be the original distribution and Q_γ the distribution under which the innovations are twisted by the candidate mean shocks over the path's $N_\gamma = \lfloor (T_{\text{conn}} + T_\rho)/\gamma \rfloor$ periods. Let A_γ be a fixed-radius tube around the candidate mean path whose terminal state lies in K_C. The Phase 2 tracking argument gives $Q_\gamma(A_\gamma) \to 1$. The logic is that this tubed path is typical under Q_γ and rare under the original P_γ. In logs, the Radon–Nikodym derivative equals $-(I_C^{\text{cand}} + o(1))/\gamma$ plus a noise term with standard deviation smaller by a $\sqrt{\gamma}$ factor. Chebyshev's inequality implies that, for every $\delta > 0$, there is an event $B_{\gamma,\delta}$ with $Q_\gamma(B_{\gamma,\delta}) \to 1$ on which the likelihood ratio $dP_\gamma/dQ_\gamma \geq \exp\{-(I_C^{\text{cand}} + \delta + o(1))/\gamma\}$. Restricting the change of measure to $A_\gamma \cap B_{\gamma,\delta}$ gives $P_\gamma(A_\gamma) \geq \exp\{-(I_C^{\text{cand}} + \delta + o(1))/\gamma\}\, Q_\gamma(A_\gamma \cap B_{\gamma,\delta})$. Both events are likely under Q_γ and δ is arbitrary, so $\liminf_{\gamma\to 0} \gamma \log P_\gamma(A_\gamma) \geq -I_C^{\text{cand}}$. The probability of the transition path under the true P_γ is given by the probability of a typical path under the "unlikely" shifted Q_γ, weighted exponentially by the action-value of the mean shocks. On the exponential scale the noise in that weighting is second-order and so the exponential cost of this route is bounded by the candidate's action value.

The final steps show that the full episode completes. Let $\mathcal{C}_\gamma$ be the event that, from a state in K_C, beliefs contract toward α_M, vigilance declines, and beliefs exit the ε_{esc}-neighborhood. Lemma 3 and Theorem 2 show $\inf_{z' \in K_C} \Pr_{z'}(\mathcal{C}_\gamma) = 1 - o(1)$. With $\mathcal{E}_\gamma$ the three-phase event and R a large constant, the Markov property for $\widetilde{Z}$ at the deterministic time N_γ gives $P_\gamma(\mathcal{E}_\gamma) \geq P_\gamma\big(A_\gamma \cap \{|x_{N_\gamma - 1}| \leq R\}\big) \inf_{z' \in K_C,\, |x| \leq R} \Pr_{(z',x)}(\mathcal{C}_\gamma)$, and the lag is bigger than R with probability below e^{-cR^2}. So $\liminf_{\gamma \to 0} \gamma \log P_\gamma(\mathcal{E}_\gamma) \geq -I_C^{\text{cand}}$.[38] In standard learning, with $V \equiv 0$, the action of the speed-ρ family is proportional to $(1+\rho)^2/\rho$ and is minimized at $\rho = 1$, consistent with Williams (2026, Proposition 3). The proposition needs one feasible speed and makes no least-cost claim for the augmented system. $\square$

Recurrence requires more than finite action. For fixed γ and i.i.d. shocks, $\widetilde{Z}_t \in \mathcal{Z}$ is a time-homogeneous Markov process on $\mathcal{Z} \times \mathbb{R}$. Suppose that it is irreducible and positive Harris recurrent, and fix γ sufficiently small. Gaussian shocks and continuity make an open subset of K_C reachable with probability bounded away from zero from every state in a compact quiescent neighborhood. The completion estimate above, then, gives exit from the ε_{esc}-neighborhood with probability bounded away from zero from every state in

[37] See the Supplementary Appendix for details. This step does not use Assumption 3.

[38] The bound is uniform over a compact neighborhood K_o of z_+ because the same twisted shocks apply to every starting state, and the tracking bound propagates a small initial displacement with a horizon-dependent constant.

K_C. Repeated returns to the quiescent neighborhood and the strong Markov property then yield infinitely many completed episodes almost surely. The two recurrence properties are assumed, not derived from the finite-cost bound.

E Robustness and applications

Proof of Proposition 4. For a monic cubic, $(c-\alpha_L)(c-\alpha_M)(c-\alpha_H) = y^3+p_0y+q_0$, where $p_0 = -\frac{1}{6}\sum_{i<j}(\alpha_i-\alpha_j)^2 < 0$, $q_0 = \prod_{j\in\{L,M,H\}}(\bar{\alpha}-\alpha_j)$, and $y \equiv c-\bar{\alpha}$. Because $G''(c) = -6y/k$, the Jensen term adds $-3Vy/k$, yielding the drift reported in the proposition. The discriminant of y^3+Ay+q_0 is $-4A^3-27q_0^2$. With $A = p_0+3V$, its zero gives V_{fold}. When the discriminant is positive, order the roots as $y_1 < y_2 < y_3$. The derivative of the monic cubic is positive at y_1, y_3 and negative at y_2. Multiplying by $-1/k$ implies the two outer mean-dynamics roots are E-stable and the middle root is unstable. When the discriminant is negative, the unique real crossing has positive monic-cubic derivative and a negative drift derivative. Finally, at the fold the equations $3y_d^2 + A = 0$ and $y_d^3 + Ay_d + q_0 = 0$ imply $2y_d^3 = q_0$, and for a monic cubic the roots sum to zero implying that the surviving root equals $-2y_d$. □

The coordination calculations follow directly from the logistic derivatives. At $p = 1/2$, $B'_\beta = \beta/2$, $B''_\beta = 0$, and $B'''_\beta = -\beta^3$. Expanding the vigilance-adjusted $B_{\beta,V}(p) = B_\beta(p)+\frac{1}{2}B''_\beta(p)V$ around $z = p-1/2$ and $v = V-V^*_{\text{logit}}$ gives $B_{\beta,V}(1/2+z)-(1/2+z) = -\frac{\beta^3}{2}vz+\frac{\beta^2(3\beta-8)}{6}z^3+O(|v||z|^3+|z|^5)$, which proves the threshold formula $V^*_{\text{logit}} = (\beta-2)/\beta^3$. The multiplicity region is $\beta > 2$, and here the cubic coefficient is negative only if $2 < \beta < 8/3$.

References

Azariadis, Costas, “The Economics of Poverty Traps Part One: Complete Markets,” *Journal of Economic Growth*, 1996, *1* (4), 449–486.

Balboni, Clare, Oriana Bandiera, Robin Burgess, Maitreesh Ghatak, and Anton Heil, “Why Do People Stay Poor?,” *Quarterly Journal of Economics*, 2022, *137* (2), 785–844.

Branch, William A., “Restricted Perceptions Equilibria and Learning in Macroeconomics,” in David Colander, ed., *Post-Walrasian Macroeconomics: Beyond the Dynamic, Stochastic General Equilibrium Model*, Cambridge University Press, 2006.

— **and Bruce McGough**, "Consistent Expectations and Misspecification in Stochastic Non-linear Economies," *Journal of Economic Dynamics and Control*, 2005, *29*, 659–676.

Branch, WIlliam A. and Bruce McGough, "Heterogeneous Expectations and Micro-foundations in Macroeconomics," *Handbook of Computational Economics, forthcoming*, 2018.

Branch, William A. and George W. Evans, "Learning about Risk and Return: A Simple Model of Bubbles and Crashes," *American Economic Journal: Macroeconomics*, 2011, *3* (3), 159–191.

— **, Bruce McGough, and Mei Zhu**, "Statistical Sunspots," *Theoretical Economics*, 2022, *17* (1), 291–329.

Bullard, James, George W. Evans, and Seppo Honkapohja, "Monetary Policy, Judgment and Near-Rational Exuberance," *American Economic Review*, 2008, *98*, 1163–1177.

Capistran, Carlos and Allan G. Timmermann, "Disagreement and Biases in Inflation Expectations," *Journal of Money, Credit, and Banking*, 2009, *41*, 365–396.

Cho, In-Koo and Jonathan Libgober, "Learning Underspecified Models," *Journal of Economic Theory*, May 2025, *226.*

— **and Kenneth Kasa**, "Learning and Model Validation," *Review of Economic Studies*, 2015, *82*, 45–82.

— **, Noah Williams, and Thomas J. Sargent**, "Escaping Nash Inflation," *Review of Economic Studies*, 2002, *69*, 1–40.

Cole, Harold L. and Timothy J. Kehoe, "Self-Fulfilling Debt Crises," *Review of Economic Studies*, 2000, *67* (1), 91–116.

Economides, Nicholas and Charles Himmelberg, "Critical Mass and Network Size with Application to the US FAX Market," Working Paper EC-95-11, New York University, Stern School of Business 1995.

Esponda, Ignacio and Demian Pouzo, "Berk–Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models," *Econometrica*, 2016, *84* (3), 1093–1130.

_ , _ , and Yuichi Yamamoto, "Asymptotic Behavior of Bayesian Learners with Misspecified Models," *Journal of Economic Theory*, 2021, *195*, 105260.

Evans, George W. and Seppo Honkapohja, *Learning and Expectations in Macroeconomics*, Princeton, New Jersey: Princeton University Press, 2001.

_ , _ , and Paul Romer, "Growth Cycles," *American Economic Review*, 1998, *88*, 495–515.

Freidlin, Mark I. and Alexander D. Wentzell, *Random Perturbations of Dynamical Systems*, 2nd ed., New York: Springer-Verlag, 1998.

Fudenberg, Drew and David K. Levine, *The Theory of Learning in Games*, Cambridge, MA: MIT Press, 1998.

_ , Giacomo Lanzani, and Philipp Strack, "Limit Points of Endogenous Misspecified Learning," *Econometrica*, 2021, *89* (3), 1065–1098.

Golubitsky, Martin and David Schaeffer, "Imperfect Bifurcation in the Presence of Symmetry," *Communications in Mathematical Physics*, 1979, *67* (3), 205–232.

Hofbauer, Josef and William H. Sandholm, "On the Global Convergence of Stochastic Fictitious Play," *Econometrica*, 2002, *70* (6), 2265–2294.

Hommes, Cars H. and Gerhard Sorger, "Consistent Expectations Equilibria," *Macroeconomic Dynamics*, 1997, *2*, 287–321.

_ and Mei Zhu, "Behavioral Learning Equilibria," *Journal of Economic Theory*, 2014, *150* (778-814).

_ , Gerhard Sorger, and Florian Wagener, "Consistency of linear forecasts in a nonlinear stochastic economy," in G.I. Bischi, C. Chiarella, and I. Sushko, eds., *Global analysis of dynamic models in economics and finance: essays in honour of Laura Gardini*, Springer, 2013.

Kasa, Kenneth, "Learning, Large Deviations, and Recurrent Currency Crises," *International Economic Review*, 2004, *45*, 141–173.

Katz, Michael L. and Carl Shapiro, "Network Externalities, Competition, and Compatibility," *American Economic Review*, 1985, *75* (3), 424–440.

Lanzani, Giacomo, "Dynamic Concern for Misspecification," *Econometrica*, 2025, *93* (4), 1333–1370.

Lybbert, Travis J., Christopher B. Barrett, Solomon Desta, and D. Layne Coppock, "Stochastic Wealth Dynamics and Risk Management among a Poor Population," *Economic Journal*, 2004, *114* (498), 750–777.

Marcet, Albert and Thomas J. Sargent, "Convergence of Least-Squares Learning in Environments with Hidden State Variables and Private Information," *Journal of Political Economy*, 1989, *97*, 1306–1322.

Santos, Paulo and Christopher B. Barrett, "Heterogeneous Wealth Dynamics: On the Roles of Risk and Ability," in Christopher B. Barrett, Michael R. Carter, and Jean-Paul Chavas, eds., *The Economics of Poverty Traps*, Chicago: University of Chicago Press, 2019, pp. 265–290.

Sargent, Thomas J., *The Conquest of American Inflation*, Princeton NJ: Princeton University Press, 1999.

Vives, Xavier, "Complementarities and Games: New Developments," *Journal of Economic Literature*, 2005, *43* (2), 437–479.

Williams, Noah, "Escape Dynamics in Learning Models," *Review of Economic Studies*, 2019, *86* (2), 882–912.

_ , "Equilibria, Folds, and Escapes," 2026. Working paper, Miami Herbert Business School.

_ , "A Learning Model of Financial Instability," *American Economic Journal: Macroeconomics*, forthcoming.