GRESHAM'S LAW OF MODEL AVERAGING

IN-KOO CHO AND KENNETH KASA

Abstract. A decision maker doubts the stationarity of his environment. In response, he uses two models, one with time-varying parameters, and another with constant parameters. Forecasts are then based on a Bayesian Model Averaging strategy, which mixes forecasts from the two models. In reality, structural parameters are constant, but the (unknown) true model features expectational feedback, which the reduced form models neglect. This feedback permits fears of parameter instability to become self-confirming. Within the context of a standard linear present value asset pricing model, we use the tools of large deviations theory to show that even though the constant parameter model would converge to the (constant parameter) Rational Expectations Equilibrium if considered in isolation, the mere presence of an unstable alternative drives it out of consideration.

JEL Classification Numbers: C63, D84

1. Introduction

2. Baseline model

2.1. Rational Expectations. Consider the following workhorse asset pricing model, in which an asset price at time t, p_t, is determined according to

$$p_t = \delta z_t + \alpha \mathsf{E}_t p_{t+1} + \sigma \epsilon_t \tag{2.1}$$

where z_t denotes observed fundamentals (e.g., dividends), and where $\alpha \in (0,1)$ is a (constant) discount rate, which determines the strength of expectational feedback. Empirically, it is close to one. The ϵ_t shock is standard Gaussian white noise. Fundamentals are assumed to evolve according to the AR(1) process

$$z_t = \rho z_{t-1} + \sigma_z \epsilon_{z,t} \tag{2.2}$$

for $\rho \in (0,1)$. The fundamentals shock, $\epsilon_{z,t}$, is standard Gaussian white noise, and is orthogonal to the price shock ϵ_t. The unique stationary rational expectations equilibrium is

$$p_t = \frac{\delta}{1-\alpha\rho} z_t + \sigma \epsilon_t. \tag{2.3}$$

Date: September 21, 2016.

We thank Mikhail Anufriev, Jaroslav Borovicka, Jim Bullard, George Evans, Lars Hansen, Seppo Honkapohja, Albert Marcet, Tom Sargent, and Noah Williams for helpful comments and discussions. This project is inspired by Evans, Honkapohja, Sargent, and Williams (2013). We are grateful to all four authors of the paper, who have allowed us to cite, compare and use their results freely. Shirley Xia provided expert research assistance. Financial support from the National Science Foundation and the Korea Research Foundation is gratefully acknowledged. We are grateful for the hospitality provided by the Federal Reserve Bank of St. Louis.

Along the equilibrium path, the dynamics of p_t can only be explained by the dynamics of fundamentals, z_t. Any excess volatility of p_t over the volatility of z_t must be soaked-up by the exogenous shock ϵ_t.

It is well known that Rational Expectations versions of this kind of model cannot explain observed asset price volatility (Shiller (1989)). We explain this volatility by assuming that agents must learn about their environment. The notion that learning might help to explain asset price volatility is hardly new (see, e.g., Timmermann (1996) for an early and influential example). However, early examples were based on least-squares learning, which exhibited asymptotic convergence to the Rational Expectations Equilibrium. This would be fine if volatility appeared to dissipate over time, but there is no evidence for this. In response, a more recent literature has assumed that agents use so-called *constant gain* learning, which discounts old data. This keeps learning alive.[1]

We allow the agent to effectively employ a time-varying gain, which is not restricted to be non-zero. We do this by supposing that agents *average* between a constant gain and a decreasing gain. Evolution of the model probability weights delivers a state-dependent gain. In some respects, our analysis resembles the gain-switching algorithm of Marcet and Nicolini (2003). However, they require the agent to commit to one or the other, whereas we permit the agent to be a Bayesian, and average between the two.

2.2. **Learning with a correct model.** Suppose an agent knows the fundamentals process in (2.2), but does not know the structural price equation in (2.1). Instead, the agent postulates the following state-space model for prices

$$p_t = \beta_t z_t + \sigma\epsilon_t \tag{2.4}$$
$$\beta_t = \beta \tag{2.5}$$

for some β. Note that the Rational Expectations equilibrium is a special case of this, with

$$\beta = \frac{\delta}{1-\alpha\rho}.$$

For now, suppose the agent adopts the dogmatic prior that parameters are constant.

$$\mathcal{M}_0: \quad \beta_t = \beta \qquad \forall t \geq 1.$$

Let $\beta_t(0)$ be the conditional mean and $\Sigma_t(0)$ be the conditional variance of the posterior belief about the unknown β. Given this belief that the true model is $\mathcal{M}_0$, $(\beta_t(0), \Sigma_t(0))$ evolves according to Kalman filter algorithm:

$$\beta_{t+1}(0) = \beta_t(0) + \left(\frac{\Sigma_t(0)}{\sigma^2 + \Sigma_t(0)z_t^2}\right) z_t(p_t - \beta_t(0)z_t) \tag{2.6}$$
$$\Sigma_{t+1}(0) = \Sigma_t(0) - \frac{(z_t\Sigma_t(0))^2}{\sigma^2 + \Sigma_t(0)z_t^2} \tag{2.7}$$

where we adopt the common assumption that β_t is based on time-$(t-1)$ information, while the time-t forecast of p_{t+1}, $\rho\beta_t z_t$, can incorporate the latest z_t observation. This

[1]For example, Benhabib and Dave (2014) show that constant gain learning can generate persistent excess volatility, and can explain why asset prices have fat-tailed distributions even when the distribution of fundamentals is thin-tailed.

assumption is made to avoid simultaneity between beliefs and observations.[2] The process, Σ_t, represents the agent's evolving estimate of the variance of β_t.

Proposition 2.1. *Given his beliefs that parameters are constant, Σ_t converges to zero at rate t^{-1}, and*

$$\beta_t \to \frac{\delta}{1-\alpha\rho}$$

with probability 1.

Proof. See Evans, Honkapohja, Sargent, and Williams (2013). □

While the agent learns the Rational Expectations equilibrium, a serious problem with $\mathcal{M}_0$ is that it fails to explain the data. Since learning is transitory, so is any learning induced parameter instability. Although there is some evidence in favor of a "Great Moderation" in the volatility of macroeconomic aggregates (at least until the recent financial crisis), there is little or no evidence for such moderation in asset markets. As a result, more recent work assumes agents view parameter instability as a permanent feature of the environment.

2.3. **Learning with a wrong model.** Now assume the agent has a different perceive law of motion, formulated by the state space model:

$$p_t = \beta_t z_t + \sigma\epsilon_t \tag{2.8}$$

$$\beta_t = \beta_{t-1} + \sigma_v v_t \tag{2.9}$$

where v_t is standard Gaussian white noise with variance σ_v^2, which is orthogonal to all other variables. In contrast to $\mathcal{M}_0$, the agent is now convinced that parameters are time-varying, which can be expressed as the parameter restriction

$$\mathcal{M}_1 : \quad \sigma_v^2 > 0.$$

A serious specification error here is that the agent does not entertain the possibility that parameters might be constant. This prevents him from ever learning the Rational Expectations equilibrium (Bullard (1992)).

The belief that $\sigma_v^2 > 0$ produces only a minor change in the Kalman filtering algorithm in (2.6) and (2.7). Let $\beta_t(1)$ and $\Sigma_t(1)$ be the mean and the variance of the posterior distribution about β_t conditioned on information at $t-1$, which is computed from a Gaussian prior. The evolution of $\beta_t(1)$ and $\Sigma_t(1)$ is dictated by the new Kalman filter:

$$\beta_{t+1}(1) = \beta_t(1) + \left(\frac{\Sigma_t(1)}{\sigma^2 + \Sigma_t(1)z_t^2}\right) z_t(p_t - \beta_t(1)z_t) \tag{2.10}$$

$$\Sigma_{t+1}(1) = \Sigma_t(1) - \frac{(z_t\Sigma_t(1))^2}{\sigma^2 + \Sigma_t(1)z_t^2} + \sigma_v^2 \tag{2.11}$$

The additional σ_v^2 term causes $\Sigma_t(1)$ to now converge to a strictly positive limit, $\bar{\Sigma} > 0$. As noted by Benveniste, Metivier, and Priouret (1990), if we assume $\sigma_v^2 \ll \sigma^2$, which we will do in what follows, we can use the approximation $\sigma^2 + \Sigma_t z_t^2 \approx \sigma^2$ in the above formulas (Σ_t is small relative to σ^2 and scales inversely with z_t^2). The Riccati equation in

[2]See Evans and Honkapohja (2001) for further discussion.

(2.11) then delivers the following approximation for the steady state variance of the state, $\bar{\Sigma} \approx \sigma \cdot \sigma_v M_z^{-1/2}$, where $M_z = \mathsf{E}(z_t^2)$ denotes the second moment of the fundamentals process. In addition, if we further assume that priors about parameter drift take the particular form, $\sigma_v^2 = \gamma\sigma^2 M_z^{-1}$, then the steady state Kalman filter takes the form of the following (discounted) recursive least-squares algorithm

$$\beta_{t+1}(1) = \beta_t(1) + \gamma M_z^{-1} z_t(p_t - \beta_t(1)z_t) \tag{2.12}$$

where priors about parameter instability are now captured by the so-called gain parameter, γ. If the agent thinks parameters are more unstable, he will use a larger gain.

An important, yet subtle, question is whether the agent's beliefs about parameter instability could become self-confirming (Sargent (2008)) due to the presence of expectational feedback. It is useful to divide this question into two parts, one related to the innovation variance, σ_v^2, and the other to the random walk nature of the dynamics.

As noted above, the innovation variance is captured by the gain parameter. Typically the gain is treated as a free parameter, and is calibrated to match some feature of the data. However, as noted by Sargent (1999), the gain in self-referential models should *not* be treated as a free parameter. It is an equilibrium object. This is because the optimal gain depends on the volatility of the data, but at the same time, the volatility of the data depends on the gain. As in a Rational Expectation Equilbrium, we have a fixed point problem. In a prescient paper, Evans and Honkapohja (1993) addressed the problem of computing this fixed point. They posed the problem as one of computing a Nash equilibrium. Under appropriate stability conditions, one can then compute the equilibrium gain by iterating on a best response mapping as usual. Later in Section 5, we exploit this idea to study the stability of our more complex Bayesian Model Averaging algorithm.

To address the second issue we need to study the dynamics of the agent's parameter estimation algorithm in (2.12). After substituting in the actual price process, (2.12) can be written as

$$\beta_{t+1}(1) = \beta_t(1) + \gamma M_z^{-1} z_t \left\{[\delta + (\alpha\rho - 1)\beta_t(1)]z_t + \sigma\epsilon_t\right\} \tag{2.13}$$

Proposition 2.2. *$\forall \sigma_v^2 > 0$, (2.13) has a stationary distribution of $\beta_t(1)$. As $\sigma_v^2 \to 0$, $\beta_t(1)$ converges weakly to the solution of the following diffusion process*

$$d\beta = (1 - \alpha\rho)\left[\frac{\delta}{1 - \alpha\rho} - \beta\right] d\tau + \sqrt{\frac{\gamma}{M_z}}\sigma dW_\tau \tag{2.14}$$

where dW_τ is the standard Wiener process.

Proof. See Kushner and Yin (1997). □

(2.14) is an Ornstein-Uhlenbeck process, which generates a stationary Gaussian distribution centered on the Rational Expectations equilibrium, $\beta = \delta/(1-\alpha\rho)$. The innovation variance is consistent with the agent's priors, since $\gamma^2\sigma^2 M_z^{-1} = \sigma_v^2$. However, $d\beta$ is autocorrelated. That is, β does *not* follow a random walk $\forall \sigma_v > 0$. Strictly speaking then, the agent's priors are misspecified.

However, traditional definitions of self-confirming equilibria presume agents have access to infinite samples. In practice, agents only have access to finite samples. Given this, we can ask whether the agent could statistically reject his prior. In the language of Hansen and Sargent (2008), we can compute the detection error probability. The detection that his prior is misspecified will be extremely difficult when the drift in (2.14) is small. This is the case when $\sigma_v > 0$ is small.

Although our agents behave as if they are Bayesian, they are in fact boundedly rational, in the sense that they cannot detect the fact that their beliefs are misspecified, based upon a finite number of samples. This is the same sense of bounded rationality in a typical model of learning in macroeconomics (e.g., Adam, Marcet, and Nicolini (2016)). We introduce an element of bounded rationality, by allowing agents' beliefs to be misspecified. However, as in Esponda and Pouzo (2015), we impose discipline by requiring beliefs to be statistically confirmed in the limit. As $\sigma_v \to 0$, the beliefs of the agents converge to a Berk-Nash equilibrium (Esponda and Pouzo (2015)), which may differ from a Nash equilibrium, as will be shown in the ensuing analysis.

3. Model averaging

Dogmatic priors (about anything) are rarely a good idea. So now suppose agents hedge their bets by entertaining the possibility that parameters are constant. Forecasts are then constructed using a traditional Bayesian Model Averaging (BMA) strategy. This strategy convexifies the model space. Let us assume that the competing models are in the mind of a single agent. This is the interpretation in Evans, Honkapohja, Sargent, and Williams (2013), which we believe a reasonable description of the behavior of a decision maker of a central bank.[3]

Let π_t denote the current probability assigned by the decision maker to $\mathcal{M}_1$ (the TVP model). Recall that $\beta_t(i)$ denote the current parameter estimate for $\mathcal{M}_i$ $\forall i \in \{0,1\}$. The policymaker's time-t forecast becomes

$$\mathsf{E}_t p_{t+1} = \rho[\pi_t \beta_t(1) + (1-\pi_t)\beta_t(0)]z_t \tag{3.15}$$

Then, the actual law of motion for p_t is

$$p_t = (\delta + \rho(\pi_t\beta_t(1) + (1-\pi_t)\beta_t(0)))\, z_t + \sigma\epsilon_t. \tag{3.16}$$

[3] An alternative way to think about model averaging is from a more decentralized perspective, where multiple agents construct and revise models, which are then marketed to a single decision maker, who does not himself construct models. This is arguably more descriptive of actual macroeconomic forecasting, and model averaging emerges quite naturally in this case. We are going to consider both possibilities, mainly for pedagogical reasons. This is because the first approach is easier to formalize, since it just involves specifying the beliefs of a single agent. In contrast, with multiple agents and multiple models, one must specify how agents perceive the forecasting efforts of other agents. If agents are aware that forecasts are being used by a policymaker, whose actions potentially influence the data-generating process, they must then form beliefs over *other* forecasters' beliefs.

It is useful to collect the formulas that dictate the evolution of the endogenous variables: $(\pi_t, \beta_t(0), \Sigma_t(0), \beta_t(1), \Sigma_t(1))$.

$$\beta_{t+1}(0) = \beta_t(0) + \left(\frac{\Sigma_t(0)}{\sigma^2 + \Sigma_t(0) z_t^2}\right) z_t(p_t - \beta_t(0) z_t) \tag{3.17}$$

$$\Sigma_{t+1}(0) = \Sigma_t(0) - \frac{(z_t \Sigma_t(0))^2}{\sigma^2 + \Sigma_t(0) z_t^2} \tag{3.18}$$

$$\beta_{t+1}(1) = \beta_t(1) + \left(\frac{\Sigma_t(1)}{\sigma^2 + \Sigma_t(1) z_t^2}\right) z_t(p_t - \beta_t(1) z_t) \tag{3.19}$$

$$\Sigma_{t+1}(1) = \Sigma_t(1) - \frac{(z_t \Sigma_t(1))^2}{\sigma^2 + \Sigma_t(1) z_t^2} + \sigma_v^2 \tag{3.20}$$

$$\frac{1}{\pi_{t+1}} - 1 = \frac{A_{t+1}(0)}{A_{t+1}(1)} \left(\frac{1}{\pi_t} - 1\right) \tag{3.21}$$

where

$$A_t(i) = \frac{1}{\sqrt{2\pi(\sigma^2 + \Sigma_t(i) z_t^2)}} \exp\left[-\frac{(p_t - \beta_t(i) z_t)^2}{2\pi(\sigma^2 + \Sigma_t(i) z_t^2)}\right] \tag{3.22}$$

$\forall i, t$.

As usual, we suppose the policymaker neglects the feedback from his forecast to the actual price process. Note that the only difference between the two parameter update equations arises from their Kalman gain which is determined by the size of $\Sigma_t(i)$. With a single forecaster who neglects feedback, these two gain sequences are independent of model averaging. Our main interest is the asymptotic properties of $(\pi_t, \beta_t(0), \beta_t(1))$ with respect to t, for a given small $\sigma_v > 0$, dictated by a system of equations (3.16)-(3.21).

Proposition 3.1. *$\forall \sigma_v > 0$,*

$$\lim_{t\to\infty} \beta_t(0) = \frac{\delta}{1 - \alpha\rho}$$

with probability 1, and $\pi_t \to \{0, 1\}$ with probability 1, as $t \to \infty$. As $t \to \infty$, the distribution of $\beta_t(1)$ converges to a stationary distribution, whose mean is $\frac{\delta}{1-\alpha\rho}$. If $\sigma_v \to 0$, the stationary distribution of $\beta_t(1)$ weakly converges to $\frac{\delta}{1-\alpha\rho}$.[4]

Proof. See Appendix B. □

[4]We use the convergence results of stochastic approximation algorithms (Kushner and Yin (1997)) and its large deviation properties (Dupuis and Kushner (1987)). All analysis requires that the stochastic processes are contained in a compact convex sets. Since we assume Gaussian shocks, $\beta_t(i)$ has a full support over $\mathbb{R}_+$. Following Kushner and Yin (1997), we assume the projection facility. $\exists B > \frac{\delta}{1-\alpha\rho}$ so that $\beta_t(i) \in [-B, B]$ for $i = 0, 1$. When $\beta_t(i) \notin [-B, B]$, we use a projection facility (Kushner and Yin (1997)) to push $\beta_t(i)$ back into $[-B, B]$. Kushner and Yin (1997) shows that the asymptotic properties of $(\pi_t, \beta_t(0), \beta_t(1))$ are not affected by the projection facility, as long as we ensure that $[-B, B]$ contains the stable point of $\beta_t(i)$. Because the Gaussian distribution has a thin tail, Dupuis and Kushner (1987) show that the large deviation properties of $(\pi_t, \beta_t(0), \beta_t(1))$ is not affected by the projection facility. For the rest of the proof, we presume that $(\pi_t, \beta_t(0), \beta_t(1)) \in [0, 1] \times [-B, B] \times [-B, B]$. To simplify notation, however, we suppress the projection facility.

The policy maker eventually learns the Rational Expectation value of coefficient $\beta = \frac{\delta}{1-\alpha\rho}$. However, we cannot conclude that the policy maker learns the Rational Expectations equilibrium, unless $\pi_t \to 0$. If $\pi_t \to 1$, then the market price p_t entails significantly more volatility than in the Rational Expectations equilibrium. Thus, a fundamental question is where π_t stays most of time between two locally stable points, 0 and 1, of π_t.

For fixed T, σ_v and $\varepsilon > 0$, define

$$T_1 = \#\{t \leq T \mid 1 - \varepsilon < \pi_t\}$$

as the number of periods during which π_t is within a small neighborhood of 1. Since 0 and 1 are the only two locally stable points, π_t stays in the neighborhood of 0 for most of the remaining $T - T_1$ periods.

Theorem 3.2. *$\forall \varepsilon > 0$,*

$$\lim_{\sigma_v \to 0} \lim_{T \to \infty} \mathsf{E} \frac{T_1}{T} = 1$$

Proof. See Appendix C. □

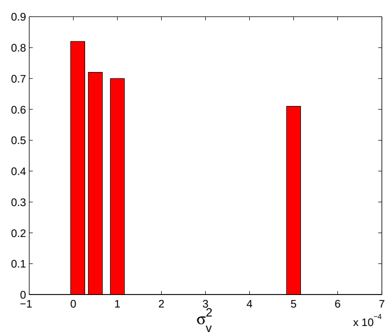

Figure 1: Probability of convergence to $\pi = 1$ after 2000 periods. As $\sigma_v \to 0$, the probability converges to 1.

The TVP model asymptotically dominates the constant parameter model, because it is better able to react to a large forecasting error. Suppose that $\pi_t \simeq 0$. By the consistency of Bayesian estimator, $\beta_t(0)$ generates only a small forecasting error with probability close to 1. However, with a vanishingly small probability, a large forecasting error occurs. In response to a large forecasting error, $\mathcal{M}_1$ adjusts $\beta_t(1)$ at a much faster rate than $\mathcal{M}_0$ adjusts $\beta_t(0)$. As a result, $\mathcal{M}_1$'s forecast starts to improve faster, and π_t increases. More importantly, as π_t increases, the forecast of $\mathcal{M}_1$ has more influence on the actual price, thus injecting more noise to the actual price process. If the data generating process becomes noisier, $\mathcal{M}_1$ responds much better than $\mathcal{M}_0$, which again increases π_t, until π_t reaches 1.

Although $\mathcal{M}_1$ is misspecified in the sense that it contains a fictitious variable without any link to fundamentals, this equilibrium must be *learned* via some adaptive process. What our result shows is that this learning process can be subverted by the mere presence of misspecified alternatives, even when the correctly specified model would converge if considered in isolation. This result therefore echoes the conclusions of Sargent (1993), who notes that adaptive learning models often need a lot of "prompting" before they converge. Elimination of misspecified alternatives can be interpreted as a form of prompting.

4. Averaging vs. Selection

Theorem 3.2 raises questions about the wisdom of using model averaging, once we entertain the possibility that models are misspecified and the data are endogenous. The fundamental problem is that model averaging forces models to *compete* with each other. The presence of TVP model can effectively change the rules of the game in its own favor, by inducing volatility that puts the constant parameters model at a competitive disadvantage.

Cho and Kasa (2015) proposed an alternative learning procedure for discriminating among multiple candidate models, which we call *model validation.* The idea behind model validation is to not *compare* models, but rather to *test* them against an externally imposed standard of statistical adequacy. If a currently employed model appears to be well specified, it continues to be used, even though some alternative model might be lurking in the background, which could statistically outperform it if given the chance. If a model is rejected, we assume that another model is randomly selected, with weights determined by historical relative performance.

In order to make the paper self-contained, let us describe the validation process of $\mathcal{M}_0$ and $\mathcal{M}_1$, following Cho and Kasa (2015). We can write the updating formula for $\beta_t(i)$ as

$$\beta_t(i) = \beta_{t-1}(i) + \eta_t(i)\Lambda_t(i) \tag{4.23}$$

$$\Lambda_t(i) = z_{t-1}[p_t - z_t\beta_{t-1}(i)] \tag{4.24}$$

where $\eta_t(i)$ is the Kalman gain of $\mathcal{M}_i$:

$$\eta_t(i) = \frac{\Sigma_t(i)}{\sigma^2 + z_{t-1}\Sigma_t(i)}.$$

Let $s_t \in \{0, 1\}$ be the model used by the policymaker, so that the actual price in period t is determined according to

$$p_t = (\delta + \rho(s_t\beta_t(1) + (1 - s_t)\beta_t(0)))\, z_t + \sigma\epsilon_t.$$

Models are tested using a recursive Lagrangean Multiplier (LM) test statistic, $\theta_t(i)$:

$$\theta_t(i) = \theta_{t-1}(i) + \eta_t^{\alpha}(i)\left[\frac{\Lambda_t^2}{\Omega_t(i)} - \theta_{t-1}^i\right] \tag{4.25}$$

$$\Omega_t(i) = \Omega_{t-1}(i) + \eta_t^{\alpha}(i)[\Lambda_t^2(i) - \Omega_{t-1}(i)] \tag{4.26}$$

where $\alpha \in (0, 1]$ is chosen to speed up the validation process. Hence, $\theta_t(i)$ is just a recursively estimated χ^2 statistic with 1 degree of freedom. We choose $\overline{\theta}$ as the test threshold. The policy maker continues to use the same model he used in period $t-1$, as long as the model passes the LM test:

$$s_t = s_{t-1}$$

if $\mathcal{M}_{s_{t-1}}$ passes the test, by satisfying $\theta_t(s_{t-1}) < \overline{\theta}$. Otherwise, $s_t = 1$ with probability one half, and $s_t = 0$ with probability one half.[5]

[5]The random selection rule here does not affect the long run distribution of s_t, as shown in Cho and Kasa (2015).

Cho and Kasa (2015) demonstrated that in the long run, the policy maker chooses a model which is the most difficult to reject. In our case, $\mathcal{M}_1$ and $\mathcal{M}_0$ are essentially the same model, except for the size of the Kalman gain. However, as $\sigma_v \to 0$, the difference of the two models vanishes, which gives an equal chance for both models to be used in the long run.

Proposition 4.1. *Define*

$$T_1^v = \#\{t \leq T \mid s_t = 1\}$$

as the number of periods when the policy maker uses $\mathcal{M}_1$ in the first T rounds, under validation dynamics. Then,

$$\lim_{\sigma_v \to 0} \lim_{T \to \infty} \mathsf{E} \frac{T_1^v}{T} = \frac{1}{2}.$$

Proof. See Cho and Kasa (2015). □

The critical difference between the model averaging dynamics and the validation dynamics is that a single model is generating the data at any point of time in the validation dynamics, while both models influence the data generating process in the model averaging process. Even in the neighborhood of $\pi_t = 0$, in which $\mathcal{M}_0$ is generating the data, any small deviation from $\pi_t = 0$ opens up the gap which $\mathcal{M}_1$ can influence the data, which in turn allows to increase π_t by injecting more noise. On the contrary, in the validation dynamics, only a significantly large the forecasting error by $\mathcal{M}_0$ can trigger to switch from $\mathcal{M}_0$ to $\mathcal{M}_1$. Since the forecasting error of $\mathcal{M}_0$ remains small most of time in the long run, $\mathcal{M}_0$ has the equal chance of being used as $\mathcal{M}_1$.

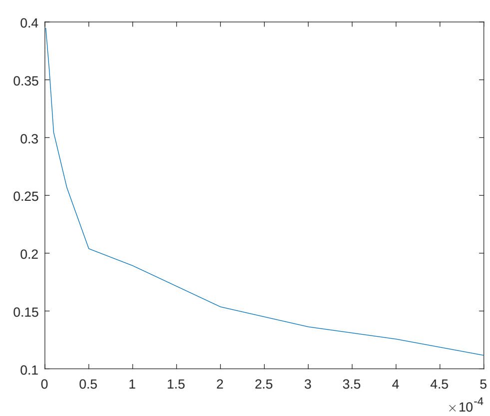

Figure 2: The horizontal axis is the values of σ_v, and the vertical axis is the proportion of time $\mathcal{M}_1$ is selected. The proportion of time is close to 0, if $\sigma_v = 0.0005$, but converges to 0.5 as $\sigma_v \to 0$.

5. Stability

Our Gresham's Law result casts doubt on the ability of agents to adaptively learn a constant parameters Rational Expectations equilibrium, unless they dogmatically believe that this is the only possible equilibrium. Here we investigate the robustness of this result to an alternative specification of the model space.

Normally, with exogenous data, it would make no difference whether a parameter known to lie in some interval is estimated by mixing between the two extremes, or by estimating it directly. With endogenous data, however, this could make a difference. What if the agent convexified the model space by estimating σ_v^2 directly, via some sort of nonlinear adaptive filtering algorithm (e.g., Mehra (1972)), or perhaps by estimating a time-varying gain instead, via an adaptive step-size algorithm (Kushner and Yang (1995))? Although $\pi = 1$ is locally stable against nonlocal alternative models, would it also be stable against local alternatives?

In this case, there is no model averaging. There is just *one* model, with σ_v^2 viewed as an unknown parameter to be estimated. To address the stability question we exploit the connection discussed in section 2.3 between σ_v^2 and the steady-state gain, γ. Because the data are endogenous, we must employ the macroeconomist's 'big K, little k' trick, which in our case we refer to as 'big Γ, little γ'. That is, our stability question can be posed as follows: Given that data are generated according to the aggregate gain parameter Γ, would an individual agent have an incentive to use a different gain, γ? If not, then $\gamma = \Gamma$ is a Nash equilibrium gain, and the associated $\sigma_v^2 > 0$ represents self-confirming parameter instability. The stability question can then be addressed by checking the (local) stability of the best response map, $\gamma = B(\Gamma)$, at the self-confirming equilibrium.

To simplify the analysis, we consider a special case, where $z_t = 1$ (i.e., $\rho = 1$ and $\sigma_z = 0$). The true model becomes

$$p_t = \delta + \alpha \mathsf{E}_t p_{t+1} + \sigma \epsilon_t \tag{5.27}$$

and the agent's perceived model becomes

$$p_t = \beta_t + \sigma \epsilon_t \tag{5.28}$$

$$\beta_t = \beta_{t-1} + \sigma_v v_t \tag{5.29}$$

where σ_v is now considered to be an unknown parameter. Note that if $\sigma_v^2 > 0$, the agent's model is misspecified. As in Sargent (1999), the agent uses a random walk to approximate a constant mean. (5.28) and (5.29) represent an example of 'random walk plus noise' model of Muth (1960), in which constant gain updating is optimal. To see this, write p_t as the following ARMA(1,1) process

$$p_t = p_{t-1} + \varepsilon_t - (1-\Gamma)\varepsilon_{t-1} \qquad \Gamma = \frac{\sqrt{4s + s^2} - s}{2} \qquad \sigma_\varepsilon^2 = \frac{\sigma^2}{1-\Gamma} \tag{5.30}$$

where $s = \sigma_v^2/\sigma^2$ is the signal-to-noise ratio. Muth (1960) showed that optimal price forecasts, $\mathsf{E}_t p_{t+1} \equiv \hat{p}_{t+1}$, evolve according to the constant gain algorithm

$$\hat{p}_{t+1} = \hat{p}_t + \Gamma(p_t - \hat{p}_t) \tag{5.31}$$

This implies that the optimal forecast of next period's price is just a geometrically distributed average of current and past prices,

$$\hat{p}_{t+1} = \left(\frac{\Gamma}{1 - (1-\Gamma)L} \right) p_t \tag{5.32}$$

Substituting this into the true model in eq. (5.27) yields the actual price process as a function of aggregate beliefs

$$\begin{aligned} p_t &= \frac{\delta}{1-\alpha} + \left(\frac{1-(1-\Gamma)L}{1-(\frac{1-\Gamma}{1-\alpha\Gamma})L}\right)\frac{\epsilon_t}{1-\alpha\Gamma} \qquad (5.33)\\ &\equiv \bar{p} + f(L;\Gamma)\tilde{\epsilon}_t \end{aligned}$$

Now for the 'big Γ, little γ' trick. Suppose prices evolve according eq. (5.33), and that an individual agent has the perceived model

$$\begin{aligned} p_t &= \frac{1-(1-\gamma)L}{1-L}u_t \qquad (5.34)\\ &\equiv h(L;\gamma)u_t \end{aligned}$$

What would be the agent's *optimal* gain? The solution of this problem defines a best response map, $\gamma = B(\Gamma)$, and a fixed point of this mapping, $\gamma = B(\gamma)$, defines a Nash equilibrium gain. Note that the agent's model is misspecified, since it omits the constant that appears in the actual prices process in eq. (5.33). The agent needs to use γ to compromise between tracking the dynamics generated by $\Gamma > 0$, and fitting the omitted constant, $\bar{p}$. This compromise is optimally resolved by minimizing the Kullback-Leibler (KLIC) distance between equations (5.33) and (5.34)[6]

$$\begin{aligned} \gamma^* = B(\Gamma) &= \text{argmin}_\gamma \left\{\mathsf{E}[h(L;\gamma)^{-1}(\bar{p} + f(L;\Gamma)\tilde{\epsilon}_t)]^2\right\} \\ &= \text{argmin}_\gamma \left\{\frac{1}{2\pi}\int_{-\pi}^{\pi}[\log H(\omega;\gamma) + \sigma_\epsilon^2 H(\omega;\gamma)^{-1}F(\omega;\Gamma) + \bar{p}^2 H(0)^{-1}]d\omega\right\} \end{aligned}$$

where $F(\omega) = f(e^{-i\omega})f(e^{i\omega})$ and $H(\omega) = h(e^{-i\omega})h(e^{i\omega})$ are the spectral densities of $f(L)$ in eq. (5.33) and $h(L)$ in eq. (5.34). Although this problem cannot be solved with pencil and paper, it is easily solved numerically. Figure 3 plots the best response map using the same benchmark parameter values as before (except, of course, $\rho = 1$ now)[7]

Not surprisingly, the agent's optimal gain increases when the external environment becomes more volatile, i.e., as Γ increases. What is more interesting is that the slope of the best response mapping is less than one. This means the equilibrium gain is *stable.* If agents believe that parameters are unstable, no single agent can do better by thinking they are less unstable. Figure 3 suggests that the best response map intersects the 45 degree line somewhere in the interval (.10, .15). This suggests that the value of σ_v^2 used for the benchmark TVP model in section 4 was a little too small, since it implied a steady-state gain of .072.

[6]See Sargent (1999, chpt. 6) for another example of this problem.

[7]Note, the unit root in the perceived model in eq. (5.34) implies that its spectral density is not well defined. (It is infinite at $\omega = 0$). In the numerical calculations, we approximate by setting $(1 - L) = (1 - \eta L)$, where $\eta = .995$. This means that our frequency domain objective is ill-equipped to find the degenerate fixed point where $\gamma = \Gamma = 0$. When this is the case, the true model exhibits i.i.d fluctuations around a mean of $\delta/(1-\alpha)$, while the agent's perceived model exhibits i.i.d fluctuations around a mean of zero. The only difference between these two processes occurs at frequency zero, which is only being approximated here.

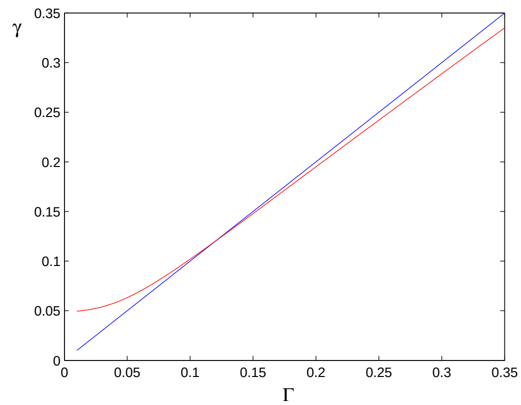

Figure 3: Best Response Mapping $\gamma = B(\Gamma)$

6. Conclusion

Parameter instability is a fact of life for applied econometricians. This paper has proposed one explanation for why this might be. We show that if econometric models are used in a less than fully understood self-referential environment, parameter instability can become a self-confirming equilibrium. Parameter estimates are unstable simply because model-builders think they *might* be unstable.

Clearly, this sort of volatility trap is an undesirable state of affairs, which raises questions about how it could be avoided. There are two main possibilities. First, not surprisingly, better theory would produce better outcomes. The agents here suffer bad outcomes because they do not fully understand their environment. If they knew the true model in eq. (2.1), they would know that data are endogenous, and would avoid reacting to their own shadows. They would simply estimate a constant parameters reduced form model. A second, and arguably more realistic possibility, is to devise econometric procedures that are more robust to misspecified endogeneity. In Cho and Kasa (2015), we argue that in this sort of environment, model selection might actually be preferable to model averaging. If agents selected either a constant or TVP model based on sequential application of a specification or hypothesis test, the constant parameter model would prevail, as it would no longer have to compete with the TVP model.

Appendix A. Preliminaries

A.1. **Dynamics of $\beta_t(0)$ and $\beta_t(1)$.** To simplify notation, let us define the Kalman gain as

$$\lambda_t(i) = \frac{\Sigma_t(i)}{\sigma^2 + \Sigma_t(i) z_t^2}$$

for $i = 1, 2$. Define $\forall \tau > 0$

$$m_i(\tau) = \inf\{K \mid \sum_{k=1}^{K} \lambda_k(i) > \tau\}$$

as the first time that $\sum_{k=1}^{K} \lambda_k(i)$ exceeds τ. Since $\lambda_k(i) > 0$ and $\sum_{k=1}^{K} \lambda_k(i) \to \infty$ with probability 1, $m_i(\tau)$ is well defined with probability 1. Similarly, define

$$\tau_K(i) = \sum_{k=1}^{K} \lambda_k(i)$$

as the size of the sum $\sum_{k=1}^{K} \lambda_k(i)$ after K rounds. $\forall K$, $\forall \tau$, consider $m(\tau_K(i) + \tau) - K$, which the number of rounds necessary for $\sum \lambda_k(i)$ to move from τ_K to $\tau_K(i) + \tau$. One can interpret $m(\tau_K(i) + \tau) - K$ as the inverse of the speed of the evolution of the associated recursive formula: if the speed of the evolution is slow, then it takes many periods to move from τ_K to $\tau_K + \tau$. We are particularly interested in the speed of evolution when K is large.

To compare the speed of evolution, we calculate

$$\lim_{K \to \infty} \frac{m(\tau_K(1) + \tau) - K}{m(\tau_K(0) + \tau) - K}.$$

If the ratio converges to 0, we say that $\beta_t(0)$ evolves at a slower time scale than $\beta_t(1)$.

Given $\sigma_v > 0$,

$$\lim_{K \to \infty} m(\tau_K(1) + \tau) - K$$

remains finite with probability 1. On the other hand,

$$\lim_{K \to \infty} m(\tau_K(0) + \tau) - K = \infty.$$

Thus, $\beta_t(0)$ evolves at a slower time scale than $\beta_t(1)$. If so, the right way to take the limit is

$$\lim_{\sigma_v \to 0} \lim_{t \to \infty}$$

because in order to move τ distance for a large K, $\beta_t(0)$ needs infinitely many more observations than $\beta_t(1)$. Based upon the order of taking limits, one can regard our exercise as calculating the long run dynamics of $(\pi_t, \beta_t(0), \beta_t(1))$ for an arbitrarily small $\sigma_v > 0$.

In order to move from $\tau_K(1)$ to $\tau_K(1) + \tau$, $\lambda_k(1)$ needs only a finite number of observations, $K_1(\tau)$. But,

$$\lim_{K \to \infty} \sum_{k=K}^{K+K_1(\tau)} \lambda_k(0) = 0.$$

As a result, $\forall \tau > 0$,

$$\lim_{K \to \infty} \beta_{K+K_1(\tau)}(0) - \beta_K(0) = 0$$

with probability 1. Therefore, in investigating the asymptotic dynamics of $\beta_t(1)$, we treat $\beta_t(0)$ as a fixed parameter. By the same token, when we investigate the asymptotic properties of $\beta_t(0)$, we assume that $\beta_t(1)$ has already reached its own stationary distribution (which is parameterized by $\beta_t(0)$).

A.2. **Dynamics of π_t.** To study the dynamics of π_t it is useful to rewrite (3.21) as follows

$$\pi_{t+1} = \pi_t + \pi_t(1-\pi_t)\left[\frac{A_{t+1}(1)/A_{t+1}(0) - 1}{1 + \pi_t(A_{t+1}(1)/A_{t+1}(0) - 1)}\right] \tag{A.35}$$

which has the familiar form of a discrete-time replicator equation, with a stochastic, state-dependent, fitness function determined by the likelihood ratio. Equation (A.35) reveals a lot about the model averaging dynamics. First, it is clear that the boundary points $\pi = \{0, 1\}$ are trivially stable fixed points, since they are absorbing. Second, we can also see that there could be an interior fixed point, where $\mathsf{E}(A_{t+1}(1)/A_{t+1}(0)) = 1$. However, we shall also see there that this fixed point is unstable. So we know already that π_t will spend most of its time near the boundary points.

Proposition A.1. *As long as the likelihoods of $\mathcal{M}_0$ and $\mathcal{M}_1$ have full support, the boundary points $\pi_t = \{0, 1\}$ are unattainable in finite time.*

Proof. With two full support probability distributions, you can never conclude that a history of any finite length couldn't have come from either of the distributions. Slightly more formally, if the distributions have full support, they are mutually absolutely continuous, so the likelihood ratio in eq. (A.35) is strictly bounded between 0 and some upper bound B. To see why $\pi_t < 1$ for all t, notice that $\pi_{t+1} < \pi_t + \pi_t(1 - \pi_t)M$ for some $M < 1$, since the likelihood ratio is bounded by B. Therefore, since $\pi + \pi(1-\pi) \in [0, 1]$ for $\pi \in [0, 1]$, we have

$$\pi_{t+1} \leq \pi_t + \pi_t(1-\pi_t)M < \pi_t + \pi_t(1-\pi_t) \leq 1$$

and so the result follows by induction. The argument for why $\pi_t > 0$ is completely symmetric. □

Since the distributions here are assumed to be Gaussian, they obviously have full support, so Proposition A.1 applies. Although the boundary points are unattainable in finite time, the replicator equation for π_t in (A.35) makes it clear that π_t will spend most of its time near these boundary points, since the relationship between π_t and π_{t+1} has the familiar logit function shape, which flattens out near the boundaries. As a result, π_t evolves very slowly near the boundary points. In fact, we shall now show that it evolves even more slowly than the t^{-1} time-scale of $\beta_t(0)$. This means that when studying the dynamics of the coefficient estimates near the boundaries, we can treat π_t as fixed.

Note that the notion of time scale is a property of a stochastic process in the right tail. That is, the time-scale measures the speed of evolution of the sample paths for large t. Although π_t can evolve faster than $\beta_t(1)$ for small t, as $t \to \infty$, we show that π_t must stay in a small neighborhood of 1 or 0, slowly converging to the limit.

Lemma A.2.

$$\mathsf{P}\left(\exists\{\pi_{t_k}\}_k, \ \textit{and } \exists \pi^* \in (0,1), \lim_{k\to\infty} \pi_{t_k} = \pi^*\right) = 0$$

and π_t evolves at a slower time scale than $\beta_t(0)$.

Proof. Fix a sequence $\{\pi_t\}$ in Π_0. Since the sequence is a subset of a compact set, it has a convergent subsequence. After renumbering the subsequence, let us assume that

$$\lim_{t\to\infty} \pi_t = \pi^* \in (0, 1)$$

since $\{\pi_t\} \in \Pi_0$. Depending upon the rate of convergence (or the time scale according to which π_t converges to π^*), we have to treat π_t has already converged to π^*.[8]

We only prove the case in which $\pi_t \to \pi^*$ according to the fastest time scale, in particular, faster than the time scale of $\beta_t(1)$. Proofs for the remaining cases follow the same logic.

Since π_t evolves according to the fastest time scale, assume that

$$\pi_t = \pi^*.$$

Since $\beta_t(1)$ evolves on a faster time scale than $\beta_t(0)$, we first let $\beta_t(1)$ reach its own "limit," and then let $\beta_t(0)$ go to its own limit point.

---

[8]If π_t evolves at a slower time scale than $\beta_t(0)$, then we fix π_t while investigating the asymptotic properties of $\beta_t(0)$. As it turns out, we obtain the same conclusion for all cases.

Fix $\sigma_v > 0$. Let $p_t^e(i)$ be the period t price forecast by model i,

$$p_t^e(1) = \beta_t(1) z_t.$$

Since

$$p_t = \alpha\rho[(1-\pi_t)\beta_t(0) + \pi_t\beta_t(1)]z_t + \delta z_t + \sigma\epsilon_t,$$

the forecast error of model 1 is

$$p_t - p_t^e(1) = [\alpha\rho(1-\pi_t)\beta_t(0) + (\alpha\rho\pi_t - 1)\beta_t(1) + \delta]\, z_t + \sigma\epsilon_t.$$

Since $\beta_t(1)$ evolves according to (2.6),

$$\lim_{t\to\infty} \mathsf{E}\left[\alpha\rho(1-\pi_t)\beta_t(0) + (\alpha\rho\pi_t - 1)\beta_t(1) + \delta\right] = 0$$

in any limit point of the Bayesian learning dynamics.[9] Define

$$\overline{\beta}(1) = \lim_{t\to 0} \mathsf{E}\beta_t(1)$$

whose value is conditioned on π_t and $\beta_t(0)$. Since

$$\lim_{t\to 0} \mathsf{E}\left[\alpha\rho(1-\pi_t)\beta_t(0) + (\alpha\rho\pi_t - 1)\overline{\beta}(1) + \delta\right] + \mathsf{E}(\alpha\rho\pi_t - 1)(\beta_t(1) - \overline{\beta}(1)) = 0.$$

Thus,

$$\overline{\beta}(1) = \frac{\alpha\rho(1-\pi_t)\beta_t(0) + \delta}{1 - \alpha\rho\pi_t} \tag{A.36}$$

for fixed $\pi_t, \beta_t(0)$. Define the deviation from the long-run mean as

$$\xi_t = \beta_t(1) - \overline{\beta}(1).$$

Model 1's mean-squared forecast error is then

$$\lim_{t\to 0} \mathsf{E}(p_t - p_t^e(1))^2 = \lim_{t\to 0} \mathsf{E} z_t^2 (\alpha\rho\pi_t - 1)^2 \sigma_\xi^2 + \sigma^2$$

Note that $\sigma_\xi^2 > 0$ if $\sigma_v > 0$, and

$$\lim_{\sigma_v^2 \to 0} \sigma_\xi^2 = 0.$$

To investigate the asymptotic properties of $\beta_t(0)$, let us write

$$\beta_t(1) = \frac{\alpha\rho(1-\pi_t)\beta_t(0) + \delta}{1 - \alpha\rho\pi_t} + \xi_t$$

Then, we can write Model 0's forecast error as

$$p_t - p_t^e(0) = z_t \left[-\frac{1-\alpha\rho}{1-\alpha\rho\pi_t} \left(\beta_t(0) - \frac{\delta}{1-\alpha\rho} \right) + \alpha\rho\pi_t\xi_t \right] + \sigma\epsilon_t.$$

Since $\beta_t(0)$ evolves according to (2.6)

$$\lim_{t\to\infty} \beta_t(0) = \frac{\delta}{1-\alpha\rho}$$

with probability 1. Thus, the mean-squared forecast error satisfies

$$\lim_{t\to\infty} \mathsf{E}(p_t - p_t^e(0))^2 = \lim_{t\to\infty} \mathsf{E} z_t^2 \sigma_\xi^2 (\alpha\rho\pi_t)^2 + \sigma^2$$

After substituting $\beta_t(0)$ into (A.36), we have

$$\lim_{\sigma_v \to 0} \lim_{t\to 0} \beta_t(1) = \frac{\delta}{1-\alpha\rho}$$

weakly. Note that

$$\lim_{t\to\infty} \frac{\mathsf{E}(p_t - p_t^e(0))^2}{\mathsf{E}(p_t - p_t^e(1))^2} > 1 \tag{A.37}$$

if and only if

$$\lim_{t\to\infty} \left(\frac{\alpha\rho\pi_t}{1-\alpha\rho\pi_t} \right)^2 > 1.$$

---

[9]Existence is implied by the tightness of the underlying space.

Now, notice that
$$\frac{\alpha\rho\pi_t}{1-\alpha\rho\pi_t} < 1$$
if and only if
$$\alpha\rho\pi_t < \frac{1}{2}.$$
Note that the left hand side is an increasing function of π_t. Hence, if (A.37) holds for some $t \geq 1$, then it holds again for $t+1$. Similarly, if (A.37) fails for some $t \geq 1$, then the same condition continues to fail for $t+1$.

Thus, π_t continues to increase or decrease, if the inequality holds in either direction. Recall that $\pi^* = \lim_{t\to\infty} \pi_t$. Convergence to π^* can occur only if (A.37) holds with equality for <u>all</u> $t \geq 1$, which is a zero probability event. We conclude that $\pi^* \in (0,1)$ occurs with probability 0. ☐

A.3. **Log odd ratio.** It is more convenient to consider the log odds ratio. Let us initialize the likelihood ratio at the prior odds ratio:
$$\frac{A_0(0)}{A_0(1)} = \frac{\pi_0(0)}{\pi_0(1)}.$$
By iteration we get
$$\frac{\pi_{t+1}(0)}{\pi_{t+1}(1)} = \frac{1}{\pi_{t+1}} - 1 = \prod_{k=0}^{t+1} \frac{A_k(0)}{A_k(1)},$$
Taking logs and dividing by $(t+1)$,
$$\frac{1}{t+1}\ln\left(\frac{1}{\pi_{t+1}} - 1\right) = \frac{1}{t+1}\sum_{k=0}^{t+1} \ln\frac{A_k(0)}{A_k(1)}.$$
Now define the average log odds ratio, ϕ_t, as follows
$$\phi_t = \frac{1}{t}\ln\left(\frac{1}{\pi_t} - 1\right) = \frac{1}{t}\ln\left(\frac{\pi_t(0)}{\pi_t(1)}\right)$$
which can be written recursively as the following stochastic approximation algorithm
$$\phi_t = \phi_{t-1} + \frac{1}{t}\left[\ln\frac{A_t(0)}{A_t(1)} - \phi_{t-1}\right].$$
Invoking well knowing results from stochastic approximation, we know that the asymptotic properties of ϕ_t are determined by the stability properties of the following ordinary differential equation (ODE)
$$\dot{\phi} = \mathsf{E}\left[\ln\frac{A_t(0)}{A_t(1)}\right] - \phi$$
which has a unique stable point
$$\phi^* = \mathsf{E}\ln\frac{A_t(0)}{A_t(1)}.$$
Note that if $\phi^* > 0$, $\pi_t \to 0$, while if $\phi^* < 0$, $\pi_t \to 1$. The focus of the ensuing analysis is to identify the conditions under which π_t converges to 1, or 0. Thus, the sign of ϕ^*, rather than its value, is an important object of investigation.

A.4. **Time scale of π_t.** Given any $\alpha \geq 1$, a simple calculation shows
$$t^\alpha(\pi_t - \pi_{t-1}) = \frac{t^\alpha(e^{(t-1)\phi_{t-1}} - e^{t\phi_t})}{(1+e^{t\phi_t})(1+e^{(t-1)\phi_{t-1}})}.$$
As $t \to \infty$, we know $\phi_t \to \phi^*$ with probability 1. Hence, we have
$$\begin{aligned}\lim_{t\to\infty} t^\alpha(\pi_t - \pi_{t-1}) &= \lim_{t\to\infty}\frac{t^\alpha\left(e^{-\phi^*}-1\right)e^{t\phi^*}}{(1+e^{t\phi^*})(1+e^{(t-1)\phi^*})} \\ &= (e^{-\phi^*}-1)\lim_{t\to\infty}\frac{t^\alpha}{(1+e^{-t\phi^*})(1+e^{t\phi^*}e^{-\phi^*})}\end{aligned}$$

Finally, notice that for both $\phi^* > 0$ and $\phi^* < 0$ the denominator converges to ∞ faster than the numerator for any $\alpha \geq 1$. Note that $\pi_t \propto \frac{1}{t}$ if

$$0 < \liminf_{t\to\infty} |t^2(\pi_t - \pi_{t-1})| \leq \limsup_{t\to\infty} |t^2(\pi_t - \pi_{t-1})| < \infty.$$

In our case, the first strict inequality is violated, which implies that π_t evolves at a rate slower than $1/t$.

A.5. **Summary.** It is helpful to summarize our findings on the time scale of three stochastic processes: π_t, $\beta_t(0)$ and $\beta_t(1)$. As indicated by (A.35), π_t evolves quickly in the interior of $[0,1]$. However, no sample path converges to $\pi^* \in (0,1)$ with a positive probability. After π_t enters a small neighborhood of $\{0,1\}$, the evolution of π_t slows down significantly. Around the neighborhood of $\{0,1\}$, we have a hierarchy of time scale among three stochastic processes. $\beta_t(1)$ evolves according to a faster time scale than $\beta_t(0)$, which evolves at a faster time scale than π_t.

Appendix B. Proof of Proposition 3.1

Although the proof follows the same logic at in the proof of Lemma A.2, we sketch the proof as a reference, but also illustrate the domain of attraction of each locally stable points, along with the description of a typical convergent path.

Fix $\sigma_v > 0$. We first investigate the properties of $(\pi_t, \beta_t(0), \beta_t(1))$ as $t \to \infty$. Since $\beta_t(1)$ evolves at the fastest time scale, we first investigate the asymptotic properties of $\beta_t(1)$ for fixed $(\pi, \beta(0))$. As we have shown in the proof of Lemma A.2, $\beta_t(1)$ has a stationary distribution, and its mean converges to

$$\overline{\beta}(1) = \frac{\alpha\rho(1-\pi_t)\beta_t(0) + \delta}{1 - \alpha\rho\pi_t}.$$

For later reference, let us define

$$\mathcal{S} = \left\{ (\pi, \beta(0), \beta(1)) \mid \beta(1) = \frac{\alpha\rho(1-\pi)\beta(0) + \delta}{1 - \alpha\rho\pi_t} \right\} \tag{B.38}$$

which is a submanifold in $\mathbb{R}^3$ as depicted in Figure 4.

Given the stationary distribution of $\beta_t(1)$, we investigate the asymptotic properties of $\beta_t(0)$, for a fixed value of π_t (in a small neighborhood of $\{0,1\}$). Again, we have shown that

$$\lim_{t\to\infty} \beta_t(0) = \frac{\delta}{1-\alpha\rho}$$

with probability 1, which implies that

$$\overline{\beta}(1) \to \frac{\delta}{1-\alpha\rho}$$

$\forall \pi_t$. Then, we observe that $\pi_t \to 1$ if and only if $\phi^* < 0$, and $\pi_t \to 0$ if and only if $\phi^* > 0$, where

$$\phi^* = \mathsf{E} \ln \frac{A_t(0)}{A_t(1)}$$

where the expectation is taken with respect to the stationary distribution in the limit as $t \to \infty$. It is convenient to consider the deterministic dynamics in terms of the time scale of $\beta_t(0)$. The domain of attraction for $(\pi, \beta(0), \beta(1)) = (0, \delta/(1-\alpha\rho), \delta/(1-\alpha\rho))$ is

$$\mathcal{D}_0 = \left\{ (\pi, \beta(0), \beta(1)) \mid \mathsf{E} \log \frac{A_t(0)}{A_t(1)} > 0 \right\}$$

where $A_t(0)$ and $A_t(1)$ are likelihood functions perceived by the agent:

$$\log A_t(1) = -\frac{[(\alpha\rho\pi_t - 1)z_t\xi_t + \sigma\epsilon_t]^2}{2(\Sigma_t(1)z_t^2 + \sigma^2)} - \frac{1}{2}\log 2\left[\Sigma_t(1)z_t^2 + \sigma^2\right]$$

and

$$\log A_t(0) = -\frac{\left[-z_t\left[\frac{1-\alpha\rho}{1-\alpha\rho\pi_t}\left(\beta_t(0) - \frac{\delta}{1-\alpha\rho}\right)\right] + \alpha\rho\pi_t z_t\xi_t + \sigma\epsilon_t\right]^2}{2(\Sigma_t(0)z_t^2 + \sigma^2)} - \frac{1}{2}\log 2\left[\Sigma_t(0)z_t^2 + \sigma^2\right].$$

It is helpful to figure out $(\pi_t, \beta_t(0))$ along the boundary of $\mathcal{D}_0$, where $\phi^* = 0$. To simplify exposition, we treat z_t as a deterministic variable, but the same analysis applied to the general case, in figuring out $(\pi_t, \beta_t(0))$ along the boundary of $\mathcal{D}_0$, where $\phi^* = 0$.

Since we are interested in the sign of ϕ^*, which is computed with respect to the probability distribution when $t \to \infty$, we substitute $\Sigma_t(1)$ by $\overline{\Sigma}$, and $\Sigma_t(0)$ by 0. After a tedious calculation (even with our simplifying assumption that z_t is deterministic), we know that

$$\phi^* = -\frac{\alpha^2\rho^2\pi_t^2(\overline{\Sigma}z_t^2)^2 + 2\sigma^2\alpha\rho\pi_t\overline{\Sigma}z_t^2}{2\sigma^2(\sigma^2 + \overline{\Sigma}z_t^2)} - \frac{z_t^2}{2\sigma^2}\left(\frac{1-\alpha\rho}{1-\alpha\rho\pi_t}\right)\left(\beta_t(0) - \frac{\delta}{1-\alpha\rho}\right)^2 + \frac{1}{2}\log\left(1 + \frac{\overline{\Sigma}z_t^2}{\sigma^2}\right).$$

Note that the right hand side is a strictly decreasing function of π_t, and $\left(\beta_t(0) - \frac{\delta}{1-\alpha\rho}\right)^2$. Thus, the contour of $(\pi_t, \beta_t(0))$ satisfying $\phi^* = 0$ is symmetric around $\beta(0) = \delta/(1-\alpha\rho)$, and $\left(\beta_t(0) - \frac{\delta}{1-\alpha\rho}\right)^2$ decreases π_t must increase, in order to satisfy $\phi^* = 0$. In particular, if $\pi_t = 0$, then

$$d(\sigma_v) = \left|\beta_t(0) - \frac{\delta}{1-\alpha\rho}\right| = \frac{\sigma}{|z_t|(1-\alpha\rho)}\sqrt{\log\left(1 + \frac{\overline{\Sigma}z_t^2}{\sigma^2}\right)} \tag{B.39}$$

which is a strictly decreasing function of $\overline{\Sigma}$, and therefore, a strictly decreasing function of σ_v. In particular,

$$\lim_{\sigma_v \to 0} d(\sigma_v) = 0.$$

Among $(\pi_t, \beta_t(0))$ satisfying $\phi^* = 0$, π_t is maximized if $\beta_t(0) = \delta/(1-\alpha\rho)$. Such π_t is the positive root of

$$\alpha^2\rho^2\overline{\Sigma}z_t^2\pi_t^2 + 2\sigma^2\alpha\rho\pi_t - \sigma^2 - \overline{\Sigma}z_t^2 = 0.$$

A simple calculation shows that if π_t is the positive root of the quadratic equation,

$$\lim_{\sigma_v \to 0} \pi_t = \frac{1}{2\alpha\rho}.$$

Thus, $\forall \epsilon > 0$, $\exists \sigma_v' > 0$ so that $\forall \sigma_v \in (0, \sigma_v')$,

$$\mathcal{D}_0 \subset \left\{(\pi, \beta(0), \beta(1)) \mid \pi \le \frac{1}{2\alpha\rho} + \epsilon\right\}.$$

Note that $\mathcal{D}_0$ looks like a pipe in $\mathbb{R}^3$, since it is independent of $\beta(1)$, as depicted in Figure 4. As $\sigma_v \to 0$, the base of $\mathcal{D}_0$ on the surface spanned by $\beta(1)$ and $\beta(0)$ shrinks, making $\mathcal{D}_0$ "thinner."

It is instructive to visualize a typical sample path of $(\pi_t, \beta_t(0), \beta_t(1))$ to a locally stable point. Suppose that $\pi_1 \in (0,1)$, and $(\pi_1, \beta_1(0), \beta_1(1))$ is outside of $\mathcal{D}_0$. Then, for a small value of t, π_t evolves rapidly toward the neighborhood of 1 or 0, whose speed of evolution may be comparable to the speed of evolution of $\beta_t(1)$, while π_t remains away from the boundary points. Since $\beta_t(1)$ evolves at the faster time scale than $\beta_t(0)$, $(\pi_t, \beta_t(0), \beta_t(1))$ evolves as if $\beta_t(0) = \beta_1(0)$, while π_t stays away from the boundary points. From the perspective of $\beta_t(0)$, $\beta_t(1)$ instantaneously moves to the neighborhood of submanifold $\mathcal{S}$, which is the reason why $\mathcal{D}_0$ is independent of $\beta_t(1)$.

$(\pi_t, \beta_t(0), \beta_t(1))$ hits the neighborhood of submanifold $\mathcal{S}$ defined by (B.38), as the distribution of $\beta_t(1)$ converges to its stationary distribution, while π_t converges to the neighborhood of either 0 or 1. Then, along the surface of $\mathcal{S}$, $(\pi_t, \beta_t(0), \beta_t(1))$ moves as $\beta_t(0)$ evolves, converging to $\frac{\delta}{1-\alpha\rho}$. After $\beta_t(0)$ reaches $\frac{\delta}{1-\alpha\rho}$ along the surface of $\mathcal{S}$ so that $\beta_t(1)$ also reaches $\frac{\delta}{1-\alpha\rho}$, π_t moves. If $(\pi_t, \beta_t(0), \beta_t(1)) \in \mathcal{S} \cap \mathcal{D}_0$, then it will converges to the limit point where $\pi_t = 0$. Otherwise, it converges to another limit point where $\pi_t = 1$.

Appendix C. Proof of Theorem 3.2

C.1. **Preliminaries.**

C.1.1. *Time scale and dynamics.* Because the three variables evolve at different times scales, the sample path in $\mathbb{R}$ has a distinctive feature. Thanks to Lemma A.2, we can assume without loss of generality, π_t is in a small neighborhood of $\{0,1\}$. First, $\beta_t(1)$ moves to submanifold $\mathcal{S}$. Second, along the surface of $\mathcal{S}$, $(\pi_t, \beta_t(0), \beta_t(1))$ moves to a locally stable point as $\beta_t(0) \to \frac{\delta}{1-\alpha\rho}$. Finally, π_t converges to 0 or 1.

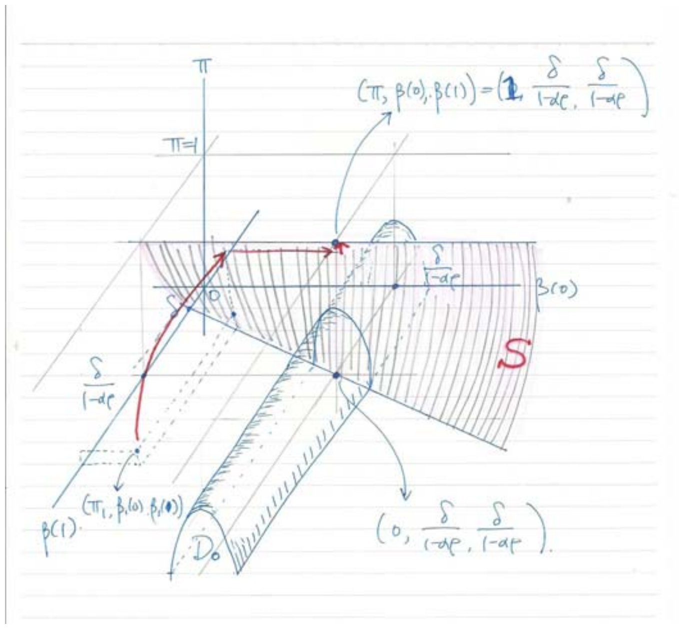

Figure 4: Red arrow shows a typical dynamics. First, $\beta_t(1)$ moves toward $\mathcal{S}$, while π_t moves to the neighborhood of 1 (or 0, depending upon the initial value). Second, $(\pi_t, \beta_t(0), \beta_t(1))$ moves along $\mathcal{S}$, as $\beta_t(0) \to \delta/(1-\alpha\rho)$. Finally, after $\beta_t(0)$ converges to $\delta/(1-\alpha\rho)$ and so does $\beta_t(1)$, π_t converges to 1.

C.1.2. *Triggering escapes.* From Figure 4, one can see that $\beta_t(1)$ does not trigger escape from $\mathcal{D}_0$ to $\mathcal{D}_1$, or $\mathcal{D}_1$ to $\mathcal{D}_0$.

It is easy to see why a large deviation by $\beta_t(1)$ from its equilibrium value $\delta/(1-\alpha\rho)$ cannot trigger $(\pi_t, \beta_t(0), \beta_t(1))$ to escape from $\mathcal{D}_0$. Note that $\mathcal{D}_0$ is independent of $\beta_t(1)$. As a result, $(\pi_t, \beta_t(0), \beta_t(1)) = (0, \frac{\delta}{1-\alpha\rho}, \beta(1)) \in \mathcal{D}_0 \ \forall \beta(1) \in \mathbb{R}$. $\mathcal{D}_0$ is the set of endogenous variables in which the mean forecasting error of $\mathcal{M}_1$ is larger than that of $\mathcal{M}_0$. The mean forecasting error is minimized when the coefficient is $\delta/(1-\alpha\rho)$. $\pi = 0$ implies that $\mathcal{M}_0$ generates a lower mean forecasting error than $\mathcal{M}_1$. If $\beta_t(1) \neq \delta/(1-\alpha\rho)$, the mean forecasting error of $\mathcal{M}_1$ increases, thus favoring $\mathcal{M}_0$, which keeps $\pi_t = 0$.

On the other hand, it is not obvious why why a large deviation by $\beta_t(1)$ from its equilibrium value $\delta/(1-\alpha\rho)$ cannot trigger $(\pi_t, \beta_t(0), \beta_t(1))$ to escape from $\mathcal{D}_1$. Note that the domain of attraction for the locally stable point where $\pi_t = 1$ is the complement of $\mathcal{D}_0$. Since

$$\mathcal{D}_0 \subset \left\{ (\pi, \beta(0), \beta(1) \mid \pi \leq \frac{1}{2\alpha\rho} \right\}.$$

Thus, the escape can occur from $\mathcal{D}_1$ only if

$$(\pi_t, \beta_t(0), \beta_t(1)) \in \left\{ (\pi, \beta(0), \beta(1) \mid \pi < \frac{1}{2\alpha\rho} \right\}.$$

That is, π_t must deviate from 1 to $\frac{1}{2\alpha\rho}$, in order to let $(\pi_t, \beta_t(0), \beta_t(1))$ escape from the domain of attraction for $(\pi, \beta(0), \beta(1)) = (1, \delta/(1-\alpha\rho), \delta/(1-\alpha\rho))$.

If π_1 at the locally stable equilibrium, $\mathcal{M}_1$ has a smaller mean forecasting error than $\mathcal{M}_0$. Suppose that $\beta_t(1)$ deviates from its equilibrium value by a large amount, which will increase the forecasting error of $\mathcal{M}_1$ and consequently, will cause π_t to decrease. However, in the neighborhood of 1, it takes an extremely large number of periods to let π_t move out of the neighborhood of 1, because π_t evolves at the slowest time scale among three stochastic processes: π_t, $\beta_t(0)$ and $\beta_t(1)$. By the time when π_t moves out of a small neighborhood of 1, the mean dynamics of $\beta_t(1)$ has already pushed back to its equilibrium value, and reduces the mean forecasting error of $\mathcal{M}_1$ below $\mathcal{M}_0$. As a result, π_t pushes back to the equilibrium value of 1.

C.2. **What to show.** Since $\beta_t(1)$ does not directly trigger the escape from one domain of attraction to another, let us focus on $(\pi, \beta(0))$, assuming that we are moving according to the time scale of $\beta_t(0)$. Note that $\mathcal{D}_0$ has a narrow symmetric pipe shape whose cross section looks like a narrow cone, centered around

$$\beta(0) = \frac{\delta}{1-\alpha\rho}$$

with the base

$$\left(\frac{\delta}{1-\alpha\rho} - d(\sigma_v), \frac{\delta}{1-\alpha\rho} + d(\sigma_v)\right)$$

along the line $\pi = 0$ where $d(\sigma_v)$ is defined as in (B.39). Recall that

$$\lim_{\sigma_v \to 0} d(\sigma_v) = 0.$$

Define

$$\bar{\pi} = \sup\{\pi \mid (\pi, \beta(0), \beta(1)) \in \mathcal{D}_0\}$$

which is $1/(2\alpha\rho)$.

Recall that

$$\phi_t = \frac{1}{t}\sum_{k=1}^{t} \log \frac{A_k(0)}{A_k(1)}.$$

Note that since $\beta_t(0), \beta_t(1) \to \frac{\delta}{1-\alpha\rho}$,

$$\phi^* = \mathsf{E} \log \frac{A_t(0)}{A_t(1)}$$

is defined for $\beta_t(0) = \beta_t(1) = \frac{\delta}{1-\alpha\rho}$, and $\pi = 1$ or 0, since $\Sigma_t(0) \to 0$ and $\Sigma_t(1) \to \overline{\Sigma}$ as $t \to 0$.

Remark C.1. *Note that as $\sigma_v \to 0$, $\overline{\Sigma} \to 0$ and consequently, $\phi^* \to 0$. One can interpret $|\phi^*|$ as the "distance" between $\mathcal{M}_0$ and $\mathcal{M}_1$. As $\sigma_v \to 0$, $\mathcal{M}_1$ and $\mathcal{M}_0$ become similar in a certain sense.*

We know that $\pi = 1$ and $\pi = 0$ are only limit points of $\{\pi_t\}$. Define ϕ^*_- as ϕ^* evaluated at $(\pi, \beta(0), \beta(1),) = (1, \frac{\delta}{1-\alpha\rho}, \frac{\delta}{1-\alpha\rho})$ and similarly, ϕ^*_+ as ϕ^* evaluated at $(\pi, \beta(0), \beta(1)) = (0, \frac{\delta}{1-\alpha\rho}, \frac{\delta}{1-\alpha\rho})$. A straightforward calculation shows

$$\phi^*_- < 0 < \phi^*_+$$

and

$$\phi^*_- + \phi^*_+ > 0.$$

For fixed $\sigma_v > 0$, define

$$r_0(\sigma_v) = -\lim_{t\to\infty} \frac{\sigma_v}{t} \log \mathsf{P}\left(\exists t,\ (\beta_t(1), \beta_t(0), \pi_t) \in \mathcal{D}_0 \mid (\beta_1(1), \beta_1(0), \pi_1) = \left(\frac{\delta}{1-\alpha\rho}, \frac{\delta}{1-\alpha\rho}, 0\right)\right)$$

and

$$r_1(\sigma_v) = -\lim_{t\to\infty} \frac{\sigma_v}{t} \log \mathsf{P}\left(\exists t,\ (\beta_t(1), \beta_t(0), \pi_t) \in \mathcal{D}_1 \mid (\beta_1(1), \beta_1(0), \pi_1) = \left(\frac{\delta}{1-\alpha\rho}, \frac{\delta}{1-\alpha\rho}, 1\right)\right)$$

Then,

$$r_0 = \lim_{\sigma_v\to 0} r_0(\sigma_v) \quad \text{and} \quad r_1 = \lim_{\sigma_v \to 0} r_1(\sigma_v)$$

which are the rate functions that dictate how difficult it is to escape from the domain of attraction of the locally stable outcome.

Let us define $\forall i \in \{0, 1\}$,

$$\tau_i^\epsilon = \inf\left\{t \mid (\pi_t, \beta_t(0), \beta_t(1)) \notin \mathcal{N}_\epsilon(\mathcal{D}_i), (\pi_1, \beta_1(0), \beta_1(1)) = \left(i, \frac{\delta}{1-\alpha\rho}, \frac{\delta}{1-\alpha\rho}\right)\right\}$$

as the first exit time from ϵ neighborhood $\mathcal{N}_\epsilon(\mathcal{D}_i)$ of $\mathcal{D}_i$. If we show that

$$r_1 > r_0,$$

then by Dupuis and Kushner (1987) and Kushner and Yin (1997), $\forall i$,

$$\tau_i^\epsilon \sim e^{tr_i}$$

in probability. Thus, the relative duration time around $\mathcal{D}_1$ satisfies

$$\lim_{t\to\infty} \mathsf{E}\frac{\tau_1^\epsilon}{\tau_0^\epsilon + \tau_1^\epsilon} = 1,$$

from which the desired conclusion follows. Thus, it remains to show that $\exists\overline{\sigma}_v > 0$ such that

$$\inf_{\sigma_v \in (0,\overline{\sigma}_v)} r_1(\sigma_v) - r_0(\sigma_v) > 0.$$

C.3. **Escape probability from $\mathcal{D}_1$.** Consider a subset of the domain of attraction $\mathcal{D}_1$ for $\mathcal{M}_1$:

$$\mathcal{D}_1' = \{(\beta(1), \beta(0), \pi) \mid \pi > \frac{1}{2\alpha\rho}\}.$$

Since $\alpha\rho > 1/2$, $\mathcal{D}_1' \neq \emptyset$. For fixed $\sigma_v > 0$, define

$$r_1^*(\sigma_v) = -\lim_{t\to\infty} \frac{\sigma_v}{t} \log \mathsf{P}\left(\exists t,\ (\beta_t(1), \beta_t(0), \pi_t) \notin \mathcal{D}_1' \mid (\beta_1(1), \beta_1(0), \pi_1) = \left(\frac{\delta}{1-\alpha\rho}, \frac{\delta}{1-\alpha\rho}, 1\right)\right)$$

and

$$r_1^* = \lim_{\sigma_v\to 0} \inf r_1^*(\sigma_v).$$

Note that

$$\exists t,\ (\beta_t(1), \beta_t(0), \pi_t) \notin \mathcal{D}_1'$$

if and only if

$$\pi_t < \bar{\pi}$$

if and only if

$$\phi_t > 0.$$

Recall

$$-\lim_{t\to\infty} \frac{\sigma_v}{t} \log \mathsf{P}\left(\exists t,\ \phi_t > 0 \mid \phi_1 = \phi_-^*\right) = r_1^*(\sigma_v).$$

Lemma C.2.

$$r_1^* = \lim_{\sigma_v\to 0} r_1^*(\sigma_v) > 0. \tag{C.40}$$

Remark C.3. *The substance of this claim is that r_1^* cannot be equal to 0. This statement would have been trivial, if ϕ_-^* is uniformly bounded away from 0. In our case, however,*

$$\lim_{\sigma_v\to 0} \phi_-^* = 0$$

which implies $\overline{\Sigma} \to 0$.

Proof. Note that

$$\phi_t > 0$$

if and only if

$$\phi_t - \phi_-^* > -\phi_-^*$$

if and only if

$$\frac{1}{t}\sum_{k=1}^t \left[\log\frac{A_t(0)}{A_t(1)} - \mathsf{E}\log\frac{A_t(0)}{A_t(1)}\right] > -\phi_-^*$$

if and only if

$$\frac{1}{t}\sum_{k=1}^t \left[\frac{\log\frac{A_t(0)}{A_t(1)} - \mathsf{E}\log\frac{A_t(0)}{A_t(1)}}{\overline{\Sigma}}\right] > -\frac{\phi_-^*}{\overline{\Sigma}}. \tag{C.41}$$

A straightforward calculation shows

$$\lim_{\sigma_v \to 0} -\frac{\phi_-^*}{\overline{\Sigma}} = \frac{\sigma_z^2}{\sigma^2}\left(\alpha\rho - \frac{1}{2}\right) > 0$$

where σ_z^2 is the stationary variance of z_t.

Remark C.4. *As $\sigma_v \to 0$, $\phi_-^* \to 0$, which makes it easier to escape from $\mathcal{D}_1$. However, as σ_v decreases, so does the the standard deviation of*

$$\log\frac{A_t(0)}{A_t(1)} - \mathsf{E}\log\frac{A_t(0)}{A_t(1)},$$

which can be interpreted as the size of deviation per each shock decreases at the same time. As a result, the number of shocks necessary for $(\pi_t, \beta_t(0), \beta_t(1))$ to escape from $\mathcal{D}_1'$ is uniformly bounded from below. As a result, the rate function is bounded away from 0.

It is tempting to conclude that we can invoke the law of large numbers to conclude that the sample average has a finite but strictly positive rate function. However,

$$\frac{\log\frac{A_t(0)}{A_t(1)} - \mathsf{E}\log\frac{A_t(0)}{A_t(1)}}{\overline{\Sigma}}$$

is not a martingale difference. Although its mean converges to 0, we cannot invoke Cramér's theorem to show the existence of a positive rate function. Instead, we shall invoke Gärtner Ellis theorem (Dembo and Zeitouni (1998)).

We can write

$$\frac{1}{t}\sum_{k=1}^{t}\left[\frac{\log\frac{A_t(0)}{A_t(1)} - \mathsf{E}\log\frac{A_t(0)}{A_t(1)}}{\overline{\Sigma}}\right] = Z_t + Y_t$$

where

$$Z_t = \frac{1}{t}\sum_{k=1}^{t}\left[\frac{\log\frac{A_t(0)}{A_t(1)} - \mathsf{E}_t\log\frac{A_t(0)}{A_t(1)}}{\overline{\Sigma}}\right]$$

and

$$Y_t = \frac{1}{t}\sum_{k=1}^{t}\left[\frac{\mathsf{E}_t\log\frac{A_t(0)}{A_t(1)} - \mathsf{E}\log\frac{A_t(0)}{A_t(1)}}{\overline{\Sigma}}\right].$$

We claim that $\forall\lambda \in \mathbb{R}$,

$$\lim_{t\to\infty}\frac{1}{t}\log\mathsf{E}e^{t\lambda Y_t} = 0.$$

A simple calculation shows

$$\mathsf{E}_t\log\frac{A_t(0)}{A_t(1)} - \mathsf{E}\log\frac{A_t(0)}{A_t(1)} = \frac{1}{2}\log\frac{\Sigma_t(1)\sigma_{z,t}^2 + \sigma^2}{\overline{\Sigma}\sigma_z^2 + \sigma^2}$$

where $\sigma_{z,t}^2$ is the conditional variance of z_t. Since $\Sigma_t(1) \to \overline{\Sigma} > 0$, $\sigma_{z,t}^2 \to \sigma_z^2$, and $\Sigma_t(1)$ is bounded, $\exists M > 0$ such that

$$\Sigma_t(1) \le M$$

and $\forall\epsilon > 0$, $\exists T(\epsilon)$ such that $\forall t \ge T(\epsilon)$,

$$\left|\mathsf{E}_t\log\frac{A_t(0)}{A_t(1)} - \mathsf{E}\log\frac{A_t(0)}{A_t(1)}\right| \le \epsilon.$$

Thus, as $t \to \infty$,

$$\frac{1}{t}\log\mathsf{E}e^{t\lambda Y_t} \le \frac{1}{t}\log\mathsf{E}e^{t|\lambda|\epsilon} + \frac{2T(\epsilon)M}{t} = |\lambda|\epsilon + \frac{2T(\epsilon)M}{t} \to |\lambda|\epsilon.$$

Since $\epsilon > 0$ is arbitrary, we have the desired conclusion.

We conclude that the H functional (a.k.a., the logarithmic moment generating function) of

$$\frac{1}{t}\sum_{k=1}^{t}\left[\frac{\log\frac{A_t(0)}{A_t(1)} - \mathsf{E}\log\frac{A_t(0)}{A_t(1)}}{\overline{\Sigma}}\right]$$

is precisely

$$H(\lambda) = \lim_{t\to\infty} \frac{1}{t} \log \mathsf{E} e^{\lambda t Z_t}.$$

That means, the large deviation properties of the left hand side of (C.41) is the same as the large deviation properties of Z_t. Since Z_t is the sample average of a martingale difference, a standard argument from large deviation theory implies that its rate function is strictly positive for given $\sigma_v > 0$. We normalized the martingale difference by dividing each term by $\overline{\Sigma}$ so that the second moment of

$$\frac{\log \frac{A_t(0)}{A_t(1)} - \mathsf{E}_t \log \frac{A_t(0)}{A_t(1)}}{\overline{\Sigma}}$$

is uniformly bounded away from 0, even in the limit as $\sigma_v \to 0$. Hence,

$$\lim_{\sigma_v\to 0} H(\lambda)$$

does not vanish to 0, which could have happened if the second moment of the marginal difference converges to 0. By applying Gärtner Ellis Theorem, we conclude that $\exists r_1^*(\sigma_v) > 0$ such that

$$\lim_{t\to\infty} \log \mathsf{P}\left(\frac{1}{t}\sum_{k=1}^{t}\left[\frac{\log \frac{A_t(0)}{A_t(1)} - \mathsf{E}\log \frac{A_t(0)}{A_t(1)}}{\overline{\Sigma}}\right] \geq -\frac{\phi_-^*}{\overline{\Sigma}}\right) = \lim_{t\to\infty} \log \mathsf{P}\left(Z_t \geq -\frac{\phi_-^*}{\overline{\Sigma}}\right) = r_1^*(\sigma_v) \tag{C.42}$$

and

$$\liminf_{\sigma_v\to 0} r_1^*(\sigma_v) = r_1^* > 0$$

as desired. □

Since $\mathcal{D}_1' \subset \mathcal{D}_1$,

$$r_1 = \lim_{\sigma_v\to 0} r_1(\sigma_v) \geq \lim_{\sigma_v\to 0} r_1^*(\sigma_v) = r_1^* > 0.$$

C.4. **Escape probability from $\mathcal{D}_0$.** Suppose that $(\pi_t, \beta_t(0), \beta_t(1))$ is in a small neighborhood of $(0, \frac{\delta}{1-\alpha\rho}, \frac{\delta}{1-\alpha\rho})$, $\beta_t(0)$ evolves according to

$$\beta_{t+1}(0) = \beta_t(0) + \frac{\Sigma_t(0) z_t^2}{\sigma^2 + \Sigma_t(0) z_t^2}\left[p_t - \beta_t(0) z_t\right].$$

At $\pi_t = 0$, the forecasting error is

$$p_t - \beta_t(0) z_t = (1-\alpha\rho)\left[\frac{\delta}{1-\alpha\rho} - \beta_t(0)\right] z_t + \sigma\epsilon_t.$$

Note that the forecast error is independent of σ_v. Following Dupuis and Kushner (1987), we can show that $\forall d > 0$, $\exists r_0^*(d) > 0$ such that

$$\lim_{t\to\infty} -\frac{1}{t}\log \mathsf{P}\left(\left|\beta_t(0) - \frac{\delta}{1-\alpha\rho}\right| > d \mid \beta_1(0) = \frac{\delta}{1-\alpha\rho}\right) = r_0^*(d)$$

and

$$\lim_{d\to 0} r_0^*(d) = 0.$$

That is, as the neighborhood of the locally stable equilibrium shrinks, it becomes easier to escape. Set $d = d(\sigma_v)$ as defined by (B.39) so that

$$\lim_{\sigma_v\to 0} r_0^*(d(\sigma_v)) = 0.$$

In principle, an exit can occur anywhere along the boundary of $\mathcal{D}_0$. By requiring that the exit must be caused by $\beta_(0)$, we make it more difficult for an exit to occur. Thus,

$$r_0(\sigma_v) \leq r_0^*(d(\sigma_v)).$$

Thus, we can find $\overline{\sigma}_v > 0$ such that $\forall \sigma_v \in (0, \overline{\sigma}_v)$,

$$r_0^*(d(\sigma_v)) < \frac{r_1^*}{2} = \frac{1}{2}\liminf_{\sigma_v\to 0} r_1^*(\sigma_v) < r_1^*(\sigma_v).$$

Thus, $\forall \sigma_v \in (0, \overline{\sigma}_v)$,

$$r_0(\sigma_v) \le r_0^*(d(\sigma_v)) < \frac{r_1^*}{2} < r_1^* \le r_1^*(\sigma_v) \le r_1(\sigma_v).$$

from which

$$\inf_{\sigma_v \in (0, \overline{\sigma}_v)} r_1(\sigma_v) - r_0^*(\sigma_v) > 0$$

follows.

References

Adam, K., A. Marcet, and J. P. Nicolini (2016): "Stock Market volatility and Learning," *Journal of Finance*, 71(1), 33–81.

Bansal, R., and A. Yaron (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59, 1481–1509.

Benhabib, J., and C. Dave (2014): "Learning, Large Deviations and Rare Events," *Review of Economic Dynamics*, 17, 367–382.

Benveniste, A., M. Metivier, and P. Priouret (1990): *Adaptive Algorithms and Stochastic Approximations*. Springer-Verlag, Berlin.

Bullard, J. (1992): "Time-Varying Parameters and Nonconvergence to Rational Expectations under Least-Squares Learning," *Economics Letters*, 40, 159–66.

Cho, I.-K., and K. Kasa (2015): "Learning and Model Validation," *Review of Economic Studies*, 82, 45–82.

Dembo, A., and O. Zeitouni (1998): *Large Deviations Techniques and Applications*. Springer-Verlag, New York, 2nd edn.

Dupuis, P., and H. J. Kushner (1987): "Asymptotic Behavior of Constrained Stochastic Approximations vis the Theory of Large Deviations," *Probability Theory and Related Fields*, 75, 223–44.

Esponda, I., and D. Pouzo (2015): "Berk-Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models," Washington University and UC Berkeley.

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

——— (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press.

Evans, G. W., S. Honkapohja, T. J. Sargent, and N. Williams (2013): "Bayesian Model Averaging, Learning, and Model Selection," in *Macroeconomics at the Service of Public Policy*, ed. by T. J. Sargent, and J. Vilmunen. Oxford University Press.

Hansen, L. P., and T. J. Sargent (2008): *Robustness*. Princeton University Press.

Kushner, H. J., and J. Yang (1995): "Analysis of Adaptive Step-Size SA Algorithms for Parameter Tracking," *IEEE Transactions on Automatic Control*, 40, 1403–10.

Kushner, H. J., and G. G. Yin (1997): *Stochastic Approximation Algorithms and Applications*. Springer-Verlag.

Marcet, A., and J. P. Nicolini (2003): "Recurrent Hyperinflations and Learning," *American Economic Review*, 93, 1476–98.

Mehra, R. K. (1972): "Approaches to Adaptive Filtering," *IEEE Transactions of Automatic Control*, 17, 693–98.

Muth, J. F. (1960): "Optimal Properties of Exponentially Weighted Forecasts," *Journal of the American Statistical Association*, 55, 299–306.

Sargent, T. J. (1993): *Bounded Rationality in Macroeconomics*. Clarendon Press.

——— (1999): *The Conquest of American Inflation*. Princeton University Press.

——— (2008): "Evolution and Intelligent Design," *American Economic Review*, 98, 5–37.

Shiller, R. J. (1989): *Market Volatility*. MIT Press.

Timmermann, A. G. (1996): "Excessive Volatility and Predictability of Stock Prices in Autoregressive Dividend Models with Learning," *Review of Economic Studies*, 63, 523–57.

Department of Economics, University of Illinois, 1407 West Gregory, Urbana IL 61801 USA

E-mail address: inkoocho@illinois.edu

URL: https://sites.google.com/site/inkoocho

Department of Economics, Simon Fraser University, 8888 University Drive, Burnaby, B.C. Canada. V5A 1S6

E-mail address: kkasa@sfu.ca

URL: http://www.sfu.ca/~kkasa