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 asset pricing model, we use the tools of large deviations theory to show that even though the constant parameter model would converge to the 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

Economists are often accused of disagreeing with each other, and indeed, when it comes to forecasting there *is* often widespread disagreement. However, there is little disagreement about how this disagreement should be resolved. At least since Bates and Granger (1969), economists have largely agreed that forecasters should hedge their bets by *averaging*. If the unknown data-generating process is a convex combination of the different models, repeated use of model averaging ultimately leads to the correct model, and disagreement disappears

---

Date: May 22, 2017.

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, who have allowed us to use their results freely. We are especially grateful to four anonymous referees and the editor for many constructive comments and suggestions. Shirley Xia provided expert research assistance. The first author gratefully acknowledge financial support from the National Science Foundation and the Korea Research Foundation (2015S1A5A2A01011999). We are grateful for the hospitality provided by the Federal Reserve Bank of St. Louis.

(Kalai and Lehrer (1993)). It is also argued that in the short-run model averaging generates more stable and accurate forecasts.[1]

This paper sounds a note of caution about this model averaging strategy. The typical analysis of model averaging takes place from the perspective of an outside econometrician, one who is attempting to understand and forecast an exogenous dynamic system. Unfortunately, this perspective is of limited relevance to macroeconomic policymakers, whose primary interest in forecasting is to influence the economy. There are of course well known procedures for forecasting in the presence of feedback and endogeneity. That is what the Rational Expectations revolution was all about, and Lucas (1976) taught us how (in principle) to do it. However, the Rational Expectations Hypothesis also makes model averaging a moot issue, since it presumes agents have common knowledge of the correct underlying model. As a result, there is no disagreement to average out.

So instead of Lucas (1976), we follow the lead of Hansen (2014), and consider a model in which 'inside' agents, who fear model misspecification, employ the same procedures recommended and studied by 'outside' econometricians. We show that when data are endogenous and models potentially misspecify this endogeneity, all of the above advantages of model averaging may be lost. In fact, we argue that a process of model validation and *selection* can actually be preferable (Cho and Kasa (2015)).

Our model features interactions among three groups of agents - A policymaker who must forecast a future price, and two competing forecasters who construct and revise models of the price process. Disagreement centers on the stationarity of the underlying environment. One forecaster thinks the environment is stationary, and so (recursively) estimates a constant parameter model. The other forecaster thinks the environment is nonstationary, and so estimates a model with drifting parameters. The policymaker is not sure who is correct, and so following standard practice, he employs a Bayesian Model Averaging (BMA) strategy, in which price forecasts are a recursively revised probability weighted average of the two forecasts. Our main result is to show that asymptotically the weight on the time-varying parameters (TVP) model converges to one, even though the underlying structural parameters are constant. As in Gresham's Law, 'bad models drive out good models'. The TVP model is bad because it generates excess volatility.[2]

[1]Timmermann (2006) emphasizes that the case for averaging rests on both solid decision-theoretic foundations, and on a wealth of practical experience. He also notes that the benefits from averaging are robust to the precise way in which forecasts are combined.

[2]Gresham's Law is named for Sir Thomas Gresham, who was a financial adviser to Queen Elizabeth I. He is often credited for noting that 'bad money drives out good money'. As DeRoover (1949) documents, 'Gresham's Law' is a bit of a misnomer. It was certainly known before Gresham, with clear descriptions

We apply our analysis to a standard linear present value asset pricing model. We show that self-confirming parameter drift can explain observed long swings in asset prices. For reasonable parameter values we find that the unconditional variance of asset prices is nearly double its Rational Expectations value. In a sense, this is not a new result. Many others have found that so-called 'constant gain' models are useful for understanding a wide variety of dynamic economic phenomena.[3] However, several nagging questions plague this literature - Why are agents so convinced that parameters are time-varying? In terms of explaining volatility, don't constant gain models "assume the result"? What if agents' beliefs were less dogmatic, and allowed for the *possibility* that parameters were constant? Our Gresham's Law result answers these questions. It shows that constant gain learning can be a self-confirming equilibrium, even when the underlying environment is stationary.

In the existing macroeconomic learning literature (e.g., Evans and Honkapohja (2001)), the *ratio* of the average speed of evolution of two different estimators is finite. Under this condition, one can approximate the asymptotic properties of the learning dynamics using the trajectory of an ordinary differential equation (ODE).[4] In our model, however, the ratio of the speed of evolution between the estimates of the TVP model and those of the constant parameter (CP) model goes to infinity, as time goes to infinity. As a result, the conventional ODE method does not apply. To prove our result, we therefore apply *two time-scale stochastic approximation* methods (Borkar (2008)). Instead of approximating the learning dynamics by ODEs *at the same time*, we approximate the learning dynamics by *a sequence* of ODEs. In particular, we first approximate the sample path of the TVP model's estimates, for fixed values of the CP model's estimates. After calculating the stationary distribution of the TVP model's estimates for each value of the CP model's estimates, we then approximate the sample path of the CP model's estimates by another ODE, with estimates of the TVP model fixed at the mean of their stationary distribution. We first prove that the model weight assigned by the policy maker to the TVP model converges to either 1 or 0. We then prove that as the parameter innovation variance in the TVP model converges to 0, the model weight converges to 1.

The intuition for why the TVP model eventually dominates is the following - When the weight on the TVP model is close to one, the world is relatively volatile (due to feedback). This makes the CP model perform relatively poorly, since it is unable to track

by Copernicus, Oresme, and even Aristophanes. There is also debate about its empirical validity (Rolnick and Weber (1986)).

[3]Examples include: Sargent (1999), Marcet and Nicolini (2003), Kasa (2004), Chakraborty and Evans (2008), Benhabib and Dave (2014), and Adam, Marcet, and Nicolini (2016).

[4]Marcet and Sargent (1989) were the first to apply this so-called 'mean ODE' approach to the macro learning literature.

the feedback-induced time-variation in the data. Of course, the tables are turned when the weight on the TVP model is close to zero. Now the world is relatively tranquil, and the TVP model suffers from additional noise, which puts it at a disadvantage. However, as long as this noise is not too large, the TVP model can exploit its ability to respond to rare sequences of shocks that generate 'large deviations' in the estimates of the CP model. In a sense, during tranquil times, the TVP model is lying in wait, ready to pounce on large deviation events. These events provide a foothold for the TVP model, which due to feedback, allows it to regain its dominance. It is tempting to speculate whether this sort of self-confirming volatility trap could be one factor in the lingering, long-term effects of rare events like financial crises.

Although framed within a particular model, our Gresham's Law result has implications of broader significance. These implications are both positive and negative. On the positive side, our result can be interpreted as providing a justification for the use of constant gain learning. As noted earlier, constant gain learning has been successfully applied to a number of questions in economics. Perhaps most closely related to our own work, Adam, Marcet, and Nicolini (2016) show that if agents think prices have a small (latent) random walk component, then this belief can become statistically self-confirming, and can generate dynamics that explain many apparent asset pricing puzzles. However, they never ask how agents came to have this particular prior in the first place, or what would happen if their priors were less dogmatic. Our result suggests that their findings would continue to apply even if agents hedged their bets by putting some weight on the stationarity of prices.[5]

On the negative side, our result casts doubt on the ability of agents to adaptively learn their way to Rational Expectations. Learning has long been used as a defense of the Rational Expectations Hypothesis, and there are results that support this conclusion (Evans and Honkapohja (2001)). However, these results endow agents with a correctly specified model, where all that needs to be learned are coefficient values. Our result shows that learning can be subverted by the mere presence of misspecified alternatives, even when the correctly specified model would converge if considered in isolation. In many ways, this result echoes the conclusions of Sargent (1993), who noted 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.

[5]Nakov and Nuno (2015) and Malmendier and Nagel (2016) provide an alternative justification of constant gain learning. They note that if overlapping generations of finitely-lived agents each use decreasing gain algorithms that place more weight on their own individual experiences, then aggregate mean beliefs can be well approximated by a constant gain algorithm.

The remainder of the paper is organized as follows. The next section presents the model. We first study learning with only one model, and discuss the sense in which beliefs converge to self-confirming equilibria. Section 3 allows the policymaker to consider multiple models, and examines the implications of Bayesian Model Averaging. We prove that as parameter drift gets small, averaging converges to the TVP model. Section 4 plots a representative sample path using conventional parameter values. It shows that small noise parameter drift can generate significant low frequency volatility. To illustrate our convergence analysis we also report mean occupancy times as a function of the innovation variance in the TVP model. Section 5 argues that a process of model testing and *selection* produces better outcomes in the presence of misspecified endogeneity. Section 6 argues that the results are robust to alternative interpretations of the model space, which permit agents to consider local alternatives. Finally, Section 7 briefly discusses some related literature, while Section 8 offers a few concluding remarks. Formal proofs are contained in the Appendix.

2. Baseline model

Our analysis is inspired by the previous work of Evans, Honkapohja, Sargent, and Williams (2013). They study a standard cobweb model, where an agent considers two models, one with constant parameters and one with time-varying parameters. The agent employs BMA when forecasting next period's price. Using simulations, they found that if expectational feedback is sufficiently strong, the weight on the TVP model often converges to one, even though the underlying parameters are constant. They offered some insightful conjectures about why this occurs, but provided no formal analysis.

2.1. **The Model.**

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. The expectations operator, E_t, denotes the beliefs of the policymaker, which may deviate from the Rational Expectations equilibrium. 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.

2.2. **Rational Expectations.** The unique stationary Rational Expectations equilibrium is

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

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)). However, these early examples were based on least-squares learning, which exhibits 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. For example, Adam, Marcet, and Nicolini (2016) show that constant gain learning can explain most observed stock market anomalies, even with modest degrees of risk aversion, while Benhabib and Dave (2014) show that constant gain learning can explain why asset prices have fat-tailed distributions, even when the distribution of fundamentals is thin-tailed. However, this recent work presumes that agents have dogmatic priors about paramter instability. We ask whether learning induced volatility survives when beliefs are less dogmatic.

2.3. **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

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

$$\beta_t = \beta_{t-1} + \sigma_v v_t \tag{2.5}$$

with $\mathrm{cov}(\epsilon, v) = 0$. The Rational Expectations equilibrium is a special case, with

$$\sigma_v = 0 \quad \text{and } \beta = \frac{\delta}{1-\alpha\rho}.$$

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

$$\mathcal{M}_0: \quad \sigma_v = 0$$

Let $(\beta_t(0), \Sigma_t(0))$ denote the conditional mean and variance of the probability distribution about unknown constant β. Given his prior, these evolve according to the following

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} can incorporate the latest z_t observation. This assumption is made to avoid simultaneity between beliefs and observations.[6] The properties of this algorithm are well known,

Proposition 2.1. *If $\alpha\rho < 1$ and the agent believes parameters are constant, $\Sigma_t \to 0$ at rate t^{-1}, and*

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

with probability 1.

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

The good news here is that if the agent *knows* that parameters are constant, he will eventually learn what they are. Given the endogeneity of the data, this is not as trivial as it might seem. In fact, that is why we must impose the stability condition $\alpha\rho < 1$. The bad news is that this model fails to explain the data. Since learning is transitory, so is any learning induced volatility. 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.

2.4. **Learning with a wrong model.** Now assume the agent has a different dogmatic prior. Suppose he is convinced that parameters are time-varying, which can be expressed as the parameter restriction

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

Although this is a wrong model from the perspective of the (unknown) Rational Expectations equilibrium, the more serious specification error here is that the agent does not even entertain the possibility that parameters might be constant. This prevents him from ever learning the Rational Expectations equilibrium (Bullard (1992)).

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

The belief that $\sigma_v^2 > 0$ produces only a minor change in the Kalman filtering algorithm in (2.6) and (2.7). The equation for the conditional mean is exactly as before, with $\beta_t(1)$ in place of $\beta_t(0)$, while the equation for the conditional variance becomes

$$\Sigma_{t+1}(1) = \Sigma_t(1) - \frac{(z_t \Sigma_t(1))^2}{\sigma^2 + \Sigma_t(1) z_t^2} + \sigma_v^2 \tag{2.8}$$

The additional σ_v^2 term causes $\Sigma_t(1)$ to 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 (2.8) 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. In addition, if we further assume that priors about parameter drift take the particular form, $\sigma_v^2 = \gamma^2 \sigma^2 M_z^{-1}$, then the steady state Kalman filter takes the form of the following (discounted) recursive least-squares algorithm

$$\beta_{t+1}(1) = \beta_t(1) + \gamma M_z^{-1} z_t (p_t - \beta_t(1) z_t) \tag{2.9}$$

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.

Although the agent's model is misspecified, the presence of feedback raises the question of whether his model could in some sense become self-confirming (Sargent (2008)).[7] That is, to what extent are the random walk priors in (2.5) consistent with the observed behavior of the parameters in the agent's model? Would an agent have an incentive to revise his prior in light of the data that are themselves (partially) generated by those priors?

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), in self-referential models, the gain is *not* 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. Evans and Honkapohja (1993) addressed the problem of computing this fixed point. They posed the problem as one

[7]Self-confirming equilibria typically presume correctly specified models, but relax identification assumptions concerning off-equilibrium path events. Technically, we are more interested in a new equilibrium concept proposed by Esponda and Pouzo (2016), called a 'Berk-Nash' equilibrium, which permits model misspecification. See below for more discussion.

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 6, 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 (2.9). After substituting in the actual price process, (2.9) can be written as

$$\beta_{t+1}(1) = \beta_t(1) + \gamma M_z^{-1} z_t \left\{[\delta + (\alpha\rho - 1)\beta_t(1)]z_t + \sigma\epsilon_t\right\} \tag{2.10}$$

In contrast to $\beta_t(0)$, the dynamics of $\beta_t(1)$ never die out. The constant gain effectively discounts old data, and so $\beta_t(1)$ never settles down. Therefore, rather than use the Law of Large Numbers to describe its asymptotics, we must use a functional Central Limit Theorem. Letting $\tilde{\beta}(1)$ be a continuous-time interpolation of $\beta_t(1)$ we have

Proposition 2.2. *If $\alpha\rho < 1$ and $\sigma_v^2 > 0$, then as $\sigma_v^2 \to 0$, $\tilde{\beta}(1)$ converges weakly to the solution of the following diffusion process*

$$d\tilde{\beta} = (1 - \alpha\rho)\left[\frac{\delta}{1 - \alpha\rho} - \tilde{\beta}\right]d\tau + \sqrt{\frac{\gamma}{M_z}}\sigma dW_\tau \tag{2.11}$$

where dW_τ is a standard Wiener process.

Proof. See Kushner and Yin (1997). □

This is an Ornstein-Uhlenbeck process, which generates a stationary Gaussian distribution centered on the Rational Expectations equilibrium. Notice the innovation variance is consistent with the agent's prior, since $\gamma^2\sigma^2 M_z^{-1} = \sigma_v^2$.[8] However, also notice that $d\tilde{\beta}$ is autocorrelated. That is, $\tilde{\beta}$ does *not* follow a random walk. Strictly speaking then, the agent's prior is misspecified. Given an infinite sample, he would be able to reject it. In practice, however, agents only have finite samples. Following Hansen and Sargent (2008), we ask - How large must the sample be before the agent could reject his prior? It is clear from (2.11) that as long as $\alpha\rho \approx 1$ and σ_v is small (so that $\tilde{\beta} \approx \delta/(1 - \alpha\rho)$), the agent will confront the problem of rejecting a unit root null against a local alternative. These tests have notoriously low power. From a practical, finite-sample, perspective the agent's prior would be (statistically) self-confirming. Adam, Marcet, and Nicolini (2016), for example,[9] used the near rationality of constant gain learning to make a convincing case that learning

[8]Although the continuous-time process in (2.11) appears to have a different innovation variance, remember that in our Central Limit scaling each unit of continuous time, τ, corresponds to γ^{-1} units of discrete time, t.

[9]Bullard, Evans, and Honkapohja (2008) would be another example.

can explain many observed asset pricing puzzles. Still, one is left to wonder how agents became so convinced that parameters are time-varying in the first place. We now show why this *must* occur.

3. Model averaging

Dogmatic priors (about anything) are rarely a good idea. So now suppose the agent hedges his 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 π_t denote the current probability assigned to $\mathcal{M}_1$ (the TVP model). Recall that $\beta_t(i)$ is the current parameter estimate for $\mathcal{M}_i$ $\forall i \in \{0,1\}$. With averaging, the policymaker's time-t forecast becomes

$$\mathsf{E}_t p_{t+1} = \rho[\pi_t \beta_t(1) + (1-\pi_t)\beta_t(0)] z_t \tag{3.12}$$

The actual law of motion for p_t is then given by

$$p_t = (\delta + \rho(\pi_t \beta_t(1) + (1-\pi_t)\beta_t(0)))\, z_t + \sigma \epsilon_t. \tag{3.13}$$

To begin, it is useful to collect the formulas that govern the evolution of the endogenous variables: $(\pi_t, \beta_t(0), \Sigma_t(0), \beta_t(1), \Sigma_t(1))$.

$$\begin{aligned}
\beta_{t+1}(0) &= \beta_t(0) + \left(\frac{\Sigma_t(0)}{\sigma^2 + \Sigma_t(0) z_t^2}\right) z_t (p_t - \beta_t(0) z_t) && (3.14)\\
\Sigma_{t+1}(0) &= \Sigma_t(0) - \frac{(z_t \Sigma_t(0))^2}{\sigma^2 + \Sigma_t(0) z_t^2} && (3.15)\\
\beta_{t+1}(1) &= \beta_t(1) + \left(\frac{\Sigma_t(1)}{\sigma^2 + \Sigma_t(1) z_t^2}\right) z_t (p_t - \beta_t(1) z_t) && (3.16)\\
\Sigma_{t+1}(1) &= \Sigma_t(1) - \frac{(z_t \Sigma_t(1))^2}{\sigma^2 + \Sigma_t(1) z_t^2} + \sigma_v^2 && (3.17)\\
\frac{1}{\pi_{t+1}} - 1 &= \frac{A_{t+1}(0)}{A_{t+1}(1)} \left(\frac{1}{\pi_t} - 1\right) && (3.18)
\end{aligned}$$

where

$$A_t(i) = \frac{1}{\sqrt{2\pi(\sigma^2 + \Sigma_t(i) z_t^2)}} \exp\left[-\frac{(p_t - \beta_t(i) z_t)^2}{2\pi(\sigma^2 + \Sigma_t(i) z_t^2)}\right] \tag{3.19}$$

is the time-t likelihood of model-i.

As usual in the learning literature, 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 $\Sigma_t(i)$. These two gain sequences are independent of model averaging. We are interested in the asymptotic properties of $(\pi_t, \beta_t(0), \beta_t(1))$ for a given small $\sigma_v > 0$. Our first result shows that parameter estimates in both models converge to the Rational Expectations value, while the model weight spends most of its time near the boundary points $\{0, 1\}$.

Proposition 3.1. *If $\alpha\rho < 1$ then $\forall \sigma_v > 0$,*

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

with probability 1, and $\pi_t \to \{0,1\}$ with probability 1. As $t \to \infty$ for a fixed $\sigma_v > 0$, the distribution of $\beta_t(1)$ converges to a stationary distribution, with mean $\frac{\delta}{1-\alpha\rho}$. If $\sigma_v \to 0$, the stationary distribution of $\beta_t(1)$ converges weakly to $\frac{\delta}{1-\alpha\rho}$.

Proof. See Appendix B. □

Convergence of the parameter estimates to the Rational Expectations value is not too surprising, since both models are forecasting the same thing using the same variable. More interesting is the result that π_t spends nearly all its time near the boundaries. For a given $\sigma_v > 0$, we cannot say ex ante which of the two locally stable limit points, $\{0, 1\}$, π_t will converge to. However, our next result shows that as $\sigma_v \to 0$, the bimodal distribution of π_t collapses to a point mass on $\pi_t = 1$.

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

$$T_1^\varepsilon = \#\{t \le T \mid \pi_t > 1 - \varepsilon\}$$

as the number of periods π_t is within a small neighborhood of 1.

Theorem 3.2. *If $\frac{1}{2} < \alpha\rho < 1$, then $\forall \varepsilon > 0$,*

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

Proof. See Appendix C. □

This is the main result of the paper. It says that the TVP model drives out the (correctly specified) constant parameter model as the drift in the TVP model gets smaller. Since the formal proof is rather involved, we relegate it to the Appendix. Here we merely attempt to provide some intuition, and offer a few interpretive comments.

First, the condition that $\alpha\rho > \frac{1}{2}$ is crucial. Without sufficiently strong feedback, the agent learns the correctly specified model, and our Gresham's Law result does not apply. Second, the order of limits is important. By letting $T \to \infty$ first, what we are doing is comparing asymptotic distributions as a parameter, σ_v, changes. The same strategy is used when using stochastic stability arguments to select among multiple Nash equilibra (Kandori, Mailath, and Rob (1993)).[10] Third, this result can be interpreted as saying that model averaging with endogenous data can exhibit a dramatic discontinuity. At precisely $\sigma_v = 0$ the two models are equivalent, and so the limiting distribution just equals the prior. However, an arbitrarily small $\sigma_v > 0$ causes the distribution to jump to a point mass on $\pi_t = 1$. Simulations reported below show that once σ_v becomes positive, the asymptotics become continuous, with the mean occupation time at $\pi_t = 1$ increasing monotonically as $\sigma_v \to 0$. This is important from a practical standpoint, since small noise random walks generate significant low frequency volatility.

The discontinuity of the asymptotic distribution of π_t with respect to σ_v^2 is driven by a fundamental asymmetric in the speed of evolution among three endogenous variables, π_t, $\beta_t(0)$ and $\beta_t(1)$), combined with the endogenous data generating process. Since the Kalman gain in (3.16) is bounded away from 0, $\beta_t(1)$ evolves at a much faster rate than $\beta_t(0)$, whose Kalman gain vanishes at the rate of $1/t$. Thus, as $t \to \infty$, $\beta_t(1)$ evolves infinitely faster than $\beta_t(0)$. π_t evolves very quickly if it is bounded away from 0 or 1, but quickly approaches, but never reaches, the boundary point of $[0, 1]$. One can show that π_t evolves at an infinitely slower speed than $\beta_t(0)$, as $t \to \infty$.

Suppose that π_t is in a small neighborhood of $\pi = 0$. Since the data is generated mostly by $\mathcal{M}_0$, the price volatility is small. $\mathcal{M}_1$ is at disadvantage, since its prediction is noisier than that of $\mathcal{M}_0$, due to the volatility of the coefficient. Since $\beta_t(0) \to \delta/(1 - \alpha\rho)$, $\beta_t(0)$ must stay within a small neighborhood of $\delta/(1 - \alpha\rho)$ with a probability close to 1. But, with a vanishingly small probability, $\beta_t(0)$ can deviate by an amount exceeding the noise in $\beta_t(1)$. In this rare event around $\pi = 0$, $\mathcal{M}_1$ generates a smaller forecasting error than $\mathcal{M}_0$, and π_t begins to increase. As π_t increases, however, the data become more volatile. As a result, the relative fitness of $\mathcal{M}_1$ increases, and π_t rapidly ascends to $\pi_t \approx 1$ in a self-reinforcing feedback loop.

In contrast, when $\pi_t \approx 1$, fluctuations in $\beta_t(1)$ do not ignite a similar self-reinforcing move toward $\pi_t \approx 0$. To see this, suppose that $\beta_t(1)$ deviates by an amount from $\delta/(1-\alpha\rho)$, and $\mathcal{M}_0$ makes a better forecast than $\mathcal{M}_1$. In a sharp contrast to the previous case, the feedback between π_t and the data does not occur, because $\beta_t(1)$ moves much faster than both π_t and $\beta_t(0)$. Before the price process can become more tranquil and put $\mathcal{M}_1$ at a further disadvantage, $\beta_1(t)$ is able to correct quickly its deviation and restore its relative fitness around $\delta/(1 - \alpha\rho)$, which keeps π_t in the neighborhood of $\pi = 1$. Escapes from

[10] Taking limits in the reverse order would be uninteresting, since the models become identical as $\sigma_v \to 0$.

$\pi_t \approx 1$ must be driven by a long sequence of 'unusual' events of π_t itself, which makes it exponentially more difficult to escape from the neighborhood of $\pi_t = 1$ than from the neighborhood of $\pi_t = 0$.

4. Simulations

We calibrate the model using parameter values that have been used in the literature, setting $\rho = 0.99$ and $\alpha = 0.96$. Since δ depends on units, we normalize it by setting $\delta = (1 - \alpha\rho)$. This implies the Rational Expectations equilibrium value, $\beta = 1.0$. In principle, the innovations variances, (σ^2, σ_z^2), could be calibrated to match those of observed assets prices and fundamentals. However, since what really matters is the comparison between actual and predicted volatility, we follow Evans, Honkapohja, Sargent, and Williams (2013) and just normalize them to unity ($\sigma^2 = \sigma_z^2 = 1$). One remaining free parameter, σ_v^2, is a crucial parameter, since it determines the agent's prior beliefs about parameter instability.

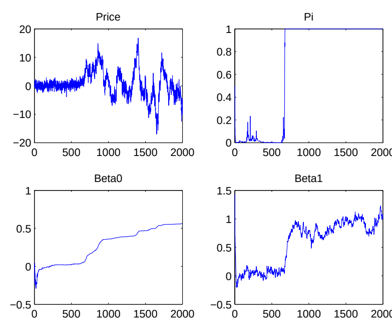

Figure 1: Sample path simulation $\sigma_v^2 = .0005$

Figure 1 plots a representative sample path with $\sigma_v^2 = .0005$. When this is the case, steady-state price volatility is 93.3% higher when $\pi = 1$ than when $\pi = 0$, which is quite significant, although less than the excess volatility detected by Shiller (1989). The higher price volatility when $\pi = 1$ is apparent. The implied steady-state gain associated with this value of σ_v^2 is $\gamma = .07$ in (2.11), which is quite typical of values used in prior empirical work. These figures also illustrate a key feature of the sample paths when σ_v^2 is relatively high: convergence to one or the other boundaries occurs relatively quickly, usually by around period 500.

For given $\sigma_v^2 > 0$, both $\pi = 0$ and $\pi = 1$ are locally stable points of the associated ordinary differential equation (ODE) of the stochastic dynamics. Let $\mathcal{D}_0$ and $\mathcal{D}_1$ be the domain of attractions for $\pi = 0$ and $\pi = 1$ of the associated ODE, respectively. The classical results of stochastic approximation (Kushner and Yin (1997)) say that if π_t visits the interior of $\mathcal{D}_i$ infinitely many times, then $\pi_t \to i$ where $i \in \{0, 1\}$, implying that the stationary distribution of π_t assigns a large weight to the neighborhood of $\pi = 0$ or $\pi = 1$. But, the same results remain silent about the probability that the expected duration of time when π_t remains in a small neighborhood of $i \in \{0, 1\}$. Our Gresham's Law result predicts that as $\sigma_v^2 \to 0$, π_t stays in a small neighborhood of $\pi = 1$ almost always, implying that the stationary distribution of π_t converges to a point mass at $\pi = 1$.

Figure 2 shows that this is indeed precisely what happens. We run $10,000$ simulations, each consisting of $T = 20,000$ time periods, initialized at $\pi \approx 1/2$ and with both parameter values near the Rational Expectations value. For each simulation we compute the mean occupancy time in a small neighborhood of each boundary (with $\varepsilon = .01$). We start with $\sigma_v^2 = 5 \times 10^{-4}$, the same value used in Figure 1, and then gradually reduce it to $\sigma_v^2 = 5 \times 10^{-6}$. When $\sigma_v^2 = 5 \times 10^{-4}$ the mean occupancy time is only about 58%. However, it rises steadily to about 80% when $\sigma_v^2 = 5 \times 10^{-6}$.[11]

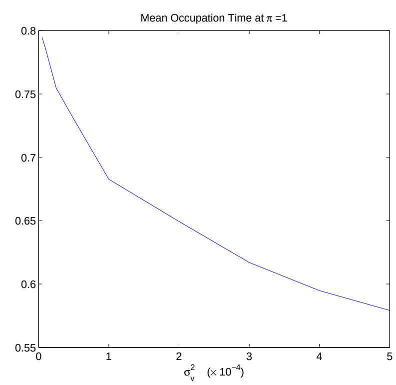

Figure 2: Occupation Times

5. 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. In general, competition is good. But here, the 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.

[11]Note, convergence here is more in the cross-sectional dimension than the time-series dimension.

Cho and Kasa (2015) proposed an alternative learning procedure for discriminating among multiple candidate models, called *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 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.

The particular validation process used by Cho and Kasa (2015) can be described as follows

$$\begin{aligned} \beta_t(i) &= \beta_{t-1}(i) + \eta_t(i)\Lambda_t(i) && (5.20) \\ \Lambda_t(i) &= z_t[p_t - z_t\beta_{t-1}(i)] && (5.21) \end{aligned}$$

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

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

Let $s_t \in \{0, 1\}$ index the model used by the policymaker in period-t, 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 Lagrange Multiplier (LM) test statistic, $\theta_t(i)$:

$$\begin{aligned} \theta_t(i) &= \theta_{t-1}(i) + \eta_t^\alpha(i)\left[\frac{\Lambda_t^2}{\Omega_t(i)} - \theta_{t-1}^i\right] && (5.22) \\ \Omega_t(i) &= \Omega_{t-1}(i) + \eta_t^\alpha(i)[\Lambda_t^2(i) - \Omega_{t-1}(i)] && (5.23) \end{aligned}$$

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.[12]

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

Cho and Kasa (2015) demonstrate that in the long run, the policy maker uses the model that is the most difficult to reject. What differentiates validation dynamics from conventional econometric testing is the fact that data are assumed to be endogenous, so each model generates data that allow it to pass the test. In our current asset pricing context, $\mathcal{M}_1$ 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 between the two models vanishes, which gives an equal chance for both models to be used in the long run.

Proposition 5.1. *Define*

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

as the number of times the policy maker uses $\mathcal{M}_1$ in the first T periods. 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 model averaging dynamics and model validation dynamics is that a *single model* is generating the data at any point of time in the validation dynamics, whereas both models influence the data with model averaging. Even in the neighborhood of $\pi_t = 0$, in which $\mathcal{M}_0$ is primarily generating the data, small increases in π_t open the door to $\mathcal{M}_1$'s influence, which in turn generates a further self-fulfilling increase in π_t. In contrast, under model validation, only a sufficiently large forecast error by $\mathcal{M}_0$ can trigger a switch from $\mathcal{M}_0$ to $\mathcal{M}_1$. Since $\mathcal{M}_0$'s forecast errors remain small most of time when $\pi \approx 0$, $\mathcal{M}_0$ retains an equal chance of being used, even in the limit as $\sigma_v^2 \to 0$.

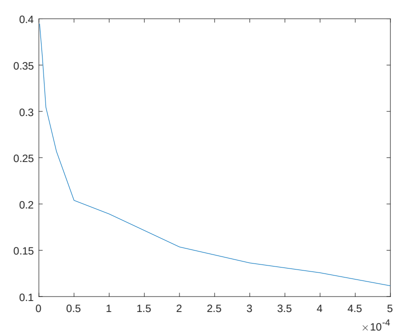

Figure 3: The horizontal axis is the values of σ_v, and the vertical axis is the proportion of time $\mathcal{M}_1$ is selected. The proportion of time is close to 0, if $\sigma_v = 0.0005$, but converges to 0.5 as $\sigma_v \to 0$.

Figure 3 verifies these predictions. We set $\overline{\theta} = 2.5$, so that it is very unlikely that single shock realizations trigger model rejections, and $\alpha = 1/2$ to speed up the validation process.

The numerical results are robust to these parameters. Figure 3 plots the mean occupancy time for $\mathcal{M}_1$ as a function of σ_v, based on 1,000 replications of 1,000 period simulations. It shows that when $\sigma_v^2 = .0005$, as in Figure 1, model validation would actually select the constant parameter model most of the time. However, as $\sigma_v^2 \to 0$, the mean occupancy converges to 0.5, as predicted.

6. 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 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

$$p_t = \delta + \alpha \mathsf{E}_t p_{t+1} + \sigma \epsilon_t \tag{6.24}$$

and the agent's perceived model becomes

$$
\begin{aligned}
p_t &= \beta_t + \sigma\epsilon_t && (6.25)\\
\beta_t &= \beta_{t-1} + \sigma_v v_t && (6.26)
\end{aligned}
$$

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. Equations (6.25) and (6.26) are an example of the '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

$$
p_t = p_{t-1} + \varepsilon_t - (1-\Gamma)\varepsilon_{t-1} \qquad \Gamma = \frac{\sqrt{4s+s^2}-s}{2} \qquad \sigma_\varepsilon^2 = \frac{\sigma^2}{1-\Gamma} \tag{6.27}
$$

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

$$
p_t^e = p_{t-1}^e + \Gamma(p_t - p_{t-1}^e) \tag{6.28}
$$

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

$$
p_t^e = \left(\frac{\Gamma}{1-(1-\Gamma)L}\right) p_t \tag{6.29}
$$

Substituting this into the true model in eq. (6.24) yields the actual price process as a function of aggregate beliefs

$$
\begin{aligned}
p_t &= \frac{\delta}{1-\alpha} + \left(\frac{1-(1-\Gamma)L}{1-(\frac{1-\Gamma}{1-\alpha\Gamma})L}\right)\frac{\epsilon_t}{1-\alpha\Gamma} && (6.30)\\
&\equiv \bar{p} + f(L;\Gamma)\tilde{\epsilon}_t
\end{aligned}
$$

Now for the 'big Γ, little γ' trick. Suppose prices evolve according eq. (6.30), and that an individual agent has the perceived model

$$
\begin{aligned}
p_t &= \frac{1-(1-\gamma)L}{1-L}u_t && (6.31)\\
&\equiv h(L;\gamma)u_t
\end{aligned}
$$

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. (6.30). 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 (6.30) and (6.31)[13]

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

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. (6.30) and $h(L)$ in eq. (6.31). Although this problem cannot be solved with pencil and paper, it is easily solved numerically. Figure 4 plots the best response map using the same benchmark parameter values as before (except, of course, $\rho = 1$ now)[14]

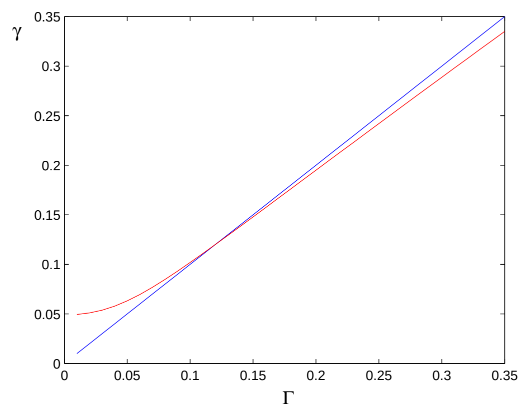

Figure 4: Best Response Mapping $\gamma = B(\Gamma)$

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 4 suggests that the best response map intersects the 45 degree

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

[14]Note, the unit root in the perceived model in eq. (6.31) 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.

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.

7. Related Literature

Although in terms of methodology our paper is most closely related to Evans, Honkapohja, Sargent, and Williams (2013), while in terms of focus it is perhaps most closely related to Adam, Marcet, and Nicolini (2016), there are a few other closely related papers, which we briefly discuss here.

7.1. **Persistent disagreement.** There are two ways to think about model averaging. One is to assume the competing models are in the mind of a single agent. Following Evans, Honkapohja, Sargent, and Williams (2013), that is how we have presented our results. However, if this is the case, one wonders why the agent does not expand the model space to include a model which nests both. In the above setting, this would mean estimating a single model where σ_v^2 is viewed as a parameter to be estimated.[15]

The second way to think about model averaging is from a more decentralized perspective, where multiple agents construct and revise models, which are then used by a single decisionmaker, who does not himself construct models. This is arguably more descriptive of actual macroeconomic forecasting, and model averaging emerges quite naturally in this case. Our stability result in the previous section suggests that no single forecaster would have an incentive to incrementally lower his σ_v^2 if other forecasters are using TVP models.

However, interpreting the model from this decentralized perspective implicitly assumes the two forecasters dogmatically stick to their own beliefs about stationarity. Is this plausible? Since they both observe the same data, one might expect from standard merging arguments that their opposing views would ultimately disappear. Although no finite amount of data could ever convince either that their beliefs about stationarity were incorrect (Kurz (1994)), complete certainty is perhaps not a reasonable standard. Could either be statistically convinced to change his views about stationarity?

Given that $\pi_t \to 1$, both forecasters end up with misspecified models. However, both will have a difficult time detecting their own misspecification. $\mathcal{M}_1$ thinks parameters exhibit a random walk, when in fact they follow a mean-reverting process centered on the Rational Expectations value. $\mathcal{M}_0$ thinks parameters are constant, but when $\pi = 1$, they instead exhibit weak mean-reverting fluctuations. Hence, $\mathcal{M}_1$ confronts the problem of

[15]Timmermann (2006) cites evidence in support of nesting, but notes that it requires agents to have access to the full information set, which is often not the case in practice.

rejecting a unit root null hypothesis. As $\sigma_v \to 0$ the power of this test converges to its size, for *any* finite sample (Blough (1992)). On the other hand, $\mathcal{M}_0$ falsely attributes $\mathcal{M}_1$'s parameter variation to his own model's error term. He therefore confronts the problem of detecting weak residual autocorrelation. Again, the power of these tests is quite low.

Hence, the two forecasters disagree about stationarity for a long time. This outcome is reminiscent of the Berk-Nash equilibrium concept proposed by Esponda and Pouzo (2016). They show that if agents' models are not econometrically identified, disagreement about model specification can persist *forever*. Although disagreement here would ultimately disappear with an infinite sample, as $\sigma_v \to 0$, it would take longer and longer to do so. In the meantime, the freedom afforded by slow parameter drift allows prices to take long excursions from their fundamental values. These excursions would not exist if the policymaker's beliefs ruled out nonstationarity from the beginning.

7.2. **Long-Run Risk.** Our previous simulations showed that self-confirming parameter drift can generate 'long swings' in asset prices. The conventional explanation of low frequency volatility in asset prices follows the long-run risk model of Bansal and Yaron (2004). They note that the combination of Epstein-Zin preferences and a small-noise near-random walk component in consumption growth can generate long swings. Hansen and Sargent (2010) offered an interesting reinterpretation of their results. They note that it is enough that agents merely *suspect* that a small noise persistent component be present. A robust portfolio/learning policy will also generate long swings, even if in reality consumption growth is i.i.d. Our paper can be seen as a combination of Bansal-Yaron and Hansen-Sargent. Here an agent's suspicions that parameters might drift can cause them to drift in reality. Unlike Bansal-Yaron and Hansen-Sargent, our explanation does not have anything to do with time-varying risk premia or robustness. The agents here are risk-neutral. Instead, our explanation relies on belief dynamics. However, like them, our explanation relies on the fact that it is empirically very difficult to detect small noise random walks with finite samples.

7.3. **Near-Rational Exuberance.** Bullard, Evans, and Honkapohja (2008) note that in practice forecasters adjust their forecasts using 'judgment'. As with our TVP forecasts, they show that such judgmental adjustments can inject self-confirming volatility into the economy, which they term 'Near-Rational Exuberance'. They rely on exactly the same sort of stability argument we outlined in Section 6. Like us, their argument also relies on model misspecifications that are difficult to detect in finite samples (i.e., small neglected moving average components). The key difference between their work and ours is that exuberance equilibria are a mere *possibility* in their model, which in practice would require coordination on a purely exogenous process. In contrast, our Gresham's Law outcome is a *necessity*, and does not require agents to coordinate on an exogenous process. Moreover, they note that exuberance equilibria would not likely occur if agents were allowed to hedge their bets by averaging, whereas our results are not only immune to averaging, but actually result from it!

7.4. **Gain-Switching.** Recursively averaging between constant and decreasing gain models can be interpreted as one way to allow for a state-dependent gain. The gain-switching algorithm of Marcet and Nicolini (2003) is another way to introduce a state-dependent gain. However, in their model agents must *commit* to one or the other at each regime of the state of the economy, whereas we permit agents to be Bayesian, and average between them. More importantly, when agents switch to a decreasing gain in Marcet and Nicolini (2003), they throw away history by re-initializing the time. They do this in order to capture some notion of a 'regime change'. In contrast, the agents in our model are not worried about discrete regime changes that make history irrelevant, but instead are trying to discriminate between gradually drifting parameters and constant parameters. Hence, history remains relevant, and the time is never reinitialized.

8. 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.

References

Adam, K., A. Marcet, and J. P. Nicolini (2016): "Stock Market Volatility and Learning," *Journal of Finance*, 71, 33–82.

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

Bates, J. M., and C. W. Granger (1969): "The Combination of Forecasts," *Operations Research Quarterly*, 20, 451–68.

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.

Blough, S. R. (1992): "The Relationship Between Power and Level for Generic Unit Root Tests in Finite Samples," *Journal of Applied Econometrics*, 7, 295–308.

Borkar, V. S. (2008): *Stochastic Approximation: A Dynamical Systems Viewpoint*. Cambridge University Press.

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

Bullard, J., G. Evans, and S. Honkapohja (2008): "Monetary Policy, Judgment, and Near-Rational Exuberance," *American Economic Review*, 98, 1163–77.

Chakraborty, A., and G. W. Evans (2008): "Can Perpetual Learning Explain the Forward Premium Puzzle?," *Journal of Monetary Economics*, 55, 477–90.

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.

DeRoover, R. (1949): *Gresham: On Foreign Exchange*. Harvard University Press.

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 (2016): "Berk-Nash Equilibrium: A Framework for Modeling Agents With Misspecified Models," *Econometrica*, 84, 1093–1130.

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. (2014): "Uncertainty Outside and Inside Economic Models," *Journal of Political Economy*, 122, 945–87.

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

——— (2010): "Fragile Beliefs and the Price of Uncertainty," *Quantitative Economics*, 1, 129–62.

Kalai, E., and E. Lehrer (1993): "Rational Learning Leads to Nash Equilibrium," *Econometrica*, 61, 1019–1045.

Kandori, M., G. Mailath, and R. Rob (1993): "Learning, Mutation and Long Run Equilibria in Games," *Econometrica*, 61, 27–56.

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

Kurz, M. (1994): "On the Structure and Diversity of Rational Beliefs," *Economic Theory*, 4, 877–900.

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.

Lucas, Jr., R. E. (1976): "Econometric Policy Evaluation: A Critique," in *The Phillips Curve and Labor Markets*, ed. by K. Brunner, and A. Meltzer. Carnegie-Rochester Conf. Series on Public Policy.

Malmendier, U., and S. Nagel (2016): "Learning from Inflation Experiences," *Quarterly Journal of Economics*, 131, 53–87.

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

Marcet, A., and T. J. Sargent (1989): "Convergence of Least Squares Learning Mechanisms in Self Referential Linear Stochastic Models," *Journal of Economic Theory*, 48, 337–368.

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.

Nakov, A., and G. Nuno (2015): "Learning from Experience in the Stock Market," *Journal of Economic Dynamics and Control*, 52, 224–39.

Rolnick, A. J., and W. E. Weber (1986): "Gresham's Law or Gresham's Fallacy?," *Journal of Political Economy*, 94, 185–89.

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. (2006): "Forecast Combinations," in *Handbook of Economic Forecasting 1*, ed. by G. Elliott, C. W. Granger, and A. Timmermann. North-Holland.

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

Appendix A. Preliminaries

Here we collect some results on time-scales, and use them to establish a 3-tiered time-scale hierarchy among $(\beta_t(0), \beta_t(1), \pi_t)$.

A.1. Dynamics of $\beta_t(0)$ and $\beta_t(1)$.

To simplify notation, write the Kalman gains 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 magnitude of the sum $\sum_{k=1}^{K} \lambda_k(i)$ after K iterations. Note that $m(\tau_K(i) + \tau) - K$ is the number of iterations necessary for $\sum \lambda_k(i)$ to move from τ_K to $\tau_K(i) + \tau$. Therefore $m(\tau_K(i) + \tau) - K$ is an inverse measure of the speed of evolution of the associated recursive formula: if the evolution speed 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 this ratio converges to 0, we say that $\beta_t(0)$ evolves on a slower time-scale than $\beta_t(1)$.

Given $\sigma_v > 0$, note that

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

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 on a slower time-scale than $\beta_t(1)$. As a result, the right way to take limits 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 $\beta_t(1)$. One can therefore 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, $K_1(\tau)$. However,

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

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

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

with probability 1. Therefore, when investigating the asymptotic dynamics of $\beta_t(1)$, we can treat $\beta_t(0)$ as fixed. By the same token, when investigating the asymptotic properties of $\beta_t(0)$, we can assume that $\beta_t(1)$ has converged to its own stationary distribution (which depends on $\beta_t(0)$).

A.2. **Dynamics of π_t.** To study the dynamics of π_t it is useful to rewrite (3.18) as follows

$$\pi_{t+1} = \pi_t + \pi_t(1-\pi_t)\left[\frac{A_{t+1}(1)/A_{t+1}(0) - 1}{1 + \pi_t(A_{t+1}(1)/A_{t+1}(0) - 1)}\right] \tag{A.32}$$

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.32) 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.32) 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

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

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

Since the distributions here are 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.32) 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, near its stable limits π_t evolves very slowly. In fact, we shall 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.

Although π_t can evolve faster than $\beta_t(1)$ for small t, as $t \to \infty$, that π_t must stay in a small neighborhood of 1 or 0, slowly converging to the limit.

Lemma A.2. *Let Π be the collection of all sample paths of $\{\pi_t\}$, and define the subset*

$$\Pi_0 = \{\{\pi_t\} \mid \textit{there is no subsequence converging to } 0 \textit{ or } 1\ \}$$

We then have

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

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 treat π_t as already having converged to π^*.[16]

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 on 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 limit.

Fix $\sigma_v > 0$. Let $p_t^e(i)$ be $\mathcal{M}_i$'s period-t forecast of p_{t+1},

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

Since

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

model 1's forecast error is

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

We know,

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

in any limit point of the Bayesian learning dynamics.[17] Define

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

Note that its value is conditioned on π_t and $\beta_t(0)$. Since

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

we have

$$\overline{\beta}(1) = \frac{\alpha\rho(1-\pi_t)\beta_t(0) + \delta}{1 - \alpha\rho\pi_t} \tag{A.33}$$

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

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

[16]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.
[17]Existence is implied by the tightness of the underlying space.

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

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

We can then 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.33), we have

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

weakly. Note that

$$\lim_{t\to\infty} \frac{\mathsf{E}(p_t - p_t^e(0))^2}{\mathsf{E}(p_t - p_t^e(1))^2} > 1 \tag{A.34}$$

if and only if

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

Finally, 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.34) holds for some $t \geq 1$, then it holds again for $t+1$. Similarly, if (A.34) 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.34) 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 odds ratio.** As usual, 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, becomes the key 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-scales of our three stochastic processes: π_t, $\beta_t(0)$ and $\beta_t(1)$. As indicated by (A.32), π_t evolves quickly in the interior of $[0,1]$. However, no sample path converges to $\pi^*\in(0,1)$ with positive probability. Once π_t enters a small neighborhood of $\{0,1\}$, the evolution of π_t slows down significantly. In the neighborhood of $\{0,1\}$, we have a hierarchy of time-scales 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 as the proof of Lemma A.2, we sketch it here as a reference. Along the way, we illustrate the domain of attraction of each locally stable point, and provide a description of a typical convergent path.

We use standard convergence results from the stochastic approximation literature (Kushner and Yin (1997)), and their large deviation properties (Dupuis and Kushner (1987)). The analysis requires that all stochastic processes are contained in compact convex sets. Since we assume Gaussian shocks, $\beta_t(i)$ has full support in $\mathbb{R}$. Following Kushner and Yin (1997), we ensure compactness by imposing a projection facility, i.e., $\exists B > \frac{\delta}{1-\alpha\rho}$ such that $\beta_t(i) \in [-B, B]$ for $i = 0, 1$. When $\beta_t(i) \notin [-B, B]$, $\beta_t(i)$ is projected back into $[-B, B]$. Kushner and Yin (1997) show that the asymptotic properties of $(\pi_t, \beta_t(0), \beta_t(1))$ are not affected by the projection facility, as long as $[-B, B]$ contains the stable point of $\beta_t(i)$. Because the Gaussian distribution has thin tails, Dupuis and Kushner (1987) are able to show that the large deviation properties of $(\pi_t, \beta_t(0), \beta_t(1))$ are 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.

We first investigate the properties of $(\pi_t, \beta_t(0), \beta_t(1))$ as $t \to \infty$ for a fixed $\sigma_v > 0$. 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, define

$$\mathcal{S} = \left\{ (\pi, \beta(0), \beta(1)) \mid \beta(1) = \frac{\alpha\rho(1 - \pi)\beta(0) + \delta}{1 - \alpha\rho\pi_t} \right\} \tag{B.35}$$

which is a submanifold in $\mathbb{R}^3$.

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, 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. It is convenient to consider the deterministic dynamics on 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 the agent's perceived likelihood functions:

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

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 characterize $(\pi_t, \beta_t(0))$ along the boundary of $\mathcal{D}_0$, where $\phi^* = 0$. To simplify exposition, we treat z_t as deterministic, but the same analysis applies to the general case with minor modifications.

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

$$\phi^* = -\frac{\alpha^2\rho^2\pi_t^2(\overline{\Sigma}z_t^2)^2 + 2\sigma^2\alpha\rho\pi_t\overline{\Sigma}z_t^2}{2\sigma^2(\sigma^2 + \overline{\Sigma}z_t^2)} - \frac{z_t^2}{2\sigma^2}\left(\frac{1-\alpha\rho}{1-\alpha\rho\pi_t}\right)\left(\beta_t(0) - \frac{\delta}{1-\alpha\rho}\right)^2 + \frac{1}{2}\log\left(1 + \frac{\overline{\Sigma}z_t^2}{\sigma^2}\right). \quad \text{(B.36)}$$

Note that the right-hand side is a decreasing function of π_t. This reflects the feedback in the system. A model's relative performance decreases when the weight on it decreases. Naturally, the right-hand side also decreases with $\left(\beta_t(0) - \frac{\delta}{1-\alpha\rho}\right)^2$, i.e., in response to deviations from the self-confirmin equilibrium. Thus, the contour of $(\pi_t, \beta_t(0))$ satisfying $\phi^* = 0$ is symmetric around $\beta(0) = \delta/(1-\alpha\rho)$, and as $\left(\beta_t(0) - \frac{\delta}{1-\alpha\rho}\right)^2$ decreases π_t must increase, in order to satisfy $\phi^* = 0$. In particular, if $\pi_t = 0$, then

$$d(\sigma_v) = \left|\beta_t(0) - \frac{\delta}{1-\alpha\rho}\right| = \frac{\sigma}{|z_t|\sqrt{1-\alpha\rho}}\sqrt{\log\left(1 + \frac{\overline{\Sigma}z_t^2}{\sigma^2}\right)} \quad \text{(B.37)}$$

which is a strictly increasing function of $\overline{\Sigma}$, and therefore, a strictly increasing 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 $\beta_t(0) = \delta/(1-\alpha\rho)$. This π_t is the positive root of

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

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 $\forall\sigma_v \in (0, \sigma_v')$,

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

Note that $\mathcal{D}_0$ looks like a pipe in $\mathbb{R}^3$, since it is independent of $\beta(1)$. As $\sigma_v \to 0$, 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 a neighborhood of 1 or 0, with a speed of evolution that 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 on a 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 a neighborhood of submanifold $\mathcal{S}$. This is 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.35), as the distribution of $\beta_t(1)$ converges to its stationary distribution, while π_t converges to a 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 converge 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 on 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 that 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.

C.1.2. *Triggering escapes.* We show why a large deviation by $\beta_t(1)$ from its mean 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 for which the mean-squared forecast error of $\mathcal{M}_1$ is larger than that of $\mathcal{M}_0$. The mean-squared forecast error is minimized when the coefficient is $\delta/(1-\alpha\rho)$. $\pi \approx 0$ implies that $\mathcal{M}_0$ generates a lower mean-squared forecast error than $\mathcal{M}_1$. If $\beta_t(1) \neq \delta/(1-\alpha\rho)$, the mean-squared forecast error of $\mathcal{M}_1$ increases, thus favoring $\mathcal{M}_0$, which keeps $\pi_t \approx 0$.

On the other hand, it is not obvious why a large deviation of $\beta_t(1)$ from its mean 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

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

Escapes 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 all the way to $\frac{1}{2\alpha\rho}$ for $(\pi_t, \beta_t(0), \beta_t(1))$ to escape from $\mathcal{D}_1$.

If $(\pi_1, \beta_t(0))$ is at the locally stable equilibrium, $\mathcal{M}_1$ has a smaller mean-squared forecast error than $\mathcal{M}_0$. Suppose that $\beta_t(1)$ deviates from its equilibrium value by a large amount, which increases the forecast error of $\mathcal{M}_1$, which causes π_t to decrease. However, when $\pi_t \approx 1$, it takes an extremely large number of periods before π_t escapes from the neighborhood of 1, because π_t evolves at the slowest time-scale. By the time π_t moves out of a small neighborhood of 1, the mean dynamics of $\beta_t(1)$ push it back to its equilibrium value, which reduces the mean-squared forecast error of $\mathcal{M}_1$ back below $\mathcal{M}_0$. As a result, the downward pressure on π_t disappears.

C.2. **What to show.** Since $\beta_t(1)$ does not directly trigger escapes from one domain of attraction to another, let us focus on $(\pi, \beta(0))$, assuming that we are moving on 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 $\beta(0) = \delta/(1-\alpha\rho)$, and 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 in (B.37). 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 therefore interpret $|\phi^*|$ as the "distance" between $\mathcal{M}_0$ and $\mathcal{M}_1$. As $\sigma_v \to 0$, the distance between $\mathcal{M}_1$ and $\mathcal{M}_0$ converges to zero.*

We know that $\pi = 1$ and $\pi = 0$ are the only limit points of $\{\pi_t\}$, although they are unattainable in finite time. Let $\mathcal{N}_{1,\epsilon} = \{\pi : |\pi - 1| < \epsilon\}$ and $\mathcal{N}_{0,\epsilon} = \{\pi : |\pi| < \epsilon\}$ be neighborhoods of these two limit points. 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 a fixed $\sigma_v > 0$, define

$$r_0(\sigma_v) = -\lim_{t\to\infty} \frac{\sigma_v}{t} \log \mathsf{P}\left((\pi_t, \beta_t(0), \beta_t(1)) \notin \mathcal{D}_0 \mid (\pi_1, \beta_1(0), \beta_1(1)) = \left(\mathcal{N}_{0,\epsilon}, \frac{\delta}{1-\alpha\rho}, \frac{\delta}{1-\alpha\rho}\right)\right) \quad \text{(C.38)}$$

and

$$r_1(\sigma_v) = -\lim_{t\to\infty} \frac{\sigma_v}{t} \log \mathsf{P}\left((\pi_t, \beta_t(0), \beta_t(1)) \notin \mathcal{D}_1 \mid (\pi_1, \beta_1(0), \beta_1(1)) = \left(\mathcal{N}_{1,\epsilon}, \frac{\delta}{1-\alpha\rho}, \frac{\delta}{1-\alpha\rho}\right)\right) \quad \text{(C.39)}$$

Then let

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

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

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

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

as the first exit time from an ϵ neighborhood $\mathcal{N}_\epsilon(\mathcal{D}_i)$ of $\mathcal{D}_i$. From Dupuis and Kushner (1987) and Kushner and Yin (1997),

$$\tau_i^\epsilon \sim e^{t r_i}$$

in probability. Thus, if we can show $r_1 > r_0$, the relative duration 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 a fixed $\sigma_v > 0$, define

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

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 be 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

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

Substituting $\pi = 1$ and $\beta_t(0) = \delta/(1-\alpha\rho)$ in (B.36), and then, applying L'Hospital's rule, we have

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

where M_z 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 $\mathcal{D}_1$, since the two models are converging to each other. However, as σ_v decreases, so does the the standard deviation of*

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

Therefore, the size of a 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 now 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 must use the 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 $\sigma^2_{z,t}$ is the conditional variance of z_t. Since $\Sigma_t(1) \to \overline{\Sigma} > 0$, $\sigma^2_{z,t} \to M_z$, and $\Sigma_t(1)$ is bounded, $\exists M > 0$ such that

$$\Sigma_t(1) \leq M$$

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

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

Thus, as $t \to \infty$,

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

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

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

This implies the large deviation properties of the left-hand side of (C.41) are 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 the 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

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

as desired. □

Finally, since $\mathcal{D}_1' \subset \mathcal{D}_1$, it is easier to escape from $\mathcal{D}_1'$ than $\mathcal{D}_1$, and we have

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

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

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

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.37) 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,

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

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.

Department of Economics, University of Illinois, 1407 West Gregory, Urbana IL 61801 USA, Hanyang University and Federal Reserve Bank of St. Louis

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`