

GRESHAM'S LAW OF MODEL AVERAGING

IN-KOO CHO AND KENNETH KASA

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

JEL Classification Numbers: C63, D84

1. INTRODUCTION

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 (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's 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.

Date: May 6, 2016.

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

¹Timmermann (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.

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 are potentially misspecified, all of the above advantages of model averaging may be lost. In fact, we argue that a process of model validation and *selection* may 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, however small the prior probability weight might be assigned to TVP model by the policy maker, 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

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.

To prove our result, we exploit time-scale separation methods (Borkar (2008)). These methods were first applied to the macroeconomic learning literature by Marcet and Sargent (1989). Time-scale separation methods allow us to effectively reduce the dimensionality of our problem. Variables that operate on a relatively slow time-scale can be fixed at their current values, while variables that operate on a relatively fast time-scale can be fixed at the means of their stationary distributions. Our problem features a hierarchy of *four* time-scales. The data operate on a relatively fast calendar time-scale. Estimates of the TVP model evolve on a slower time-scale, determined by the innovation variance of its parameters. Estimates of the constant parameters model evolve even slower, on a time-scale determined by the inverse of the historical sample size. Finally, the model weight evolves on a variable time scale, but spends most of its time in the neighborhood of either 0 or 1, where it evolves on a time-scale that is even slower than that

²Gresham’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’. Not surprisingly, ‘Gresham’s Law’ is a bit of a misnomer. As DeRoover (1949) documents, it was certainly known before Gresham, with clear descriptions by Copernicus, Oresme, and even Aristophanes. There is also debate about its empirical validity (Rolnick and Weber (1986)).

³Examples include: Sargent (1999), Cho, Williams, and Sargent (2002), Marcet and Nicolini (2003), Kasa (2004), Chakraborty and Evans (2008), and Benhabib and Dave (2014).

of the constant parameters model. The fact that the TVP model evolves on a faster time-scale than the constant parameters model is crucial to our Gresham's Law result.

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 constant parameters model perform relatively poorly, since it is unable to track 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 isn't 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 constant parameters 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.

In addition to providing a warning about the potential dangers of model averaging with endogenous data, our paper also provides an example of a new equilibrium concept in learning models. Traditional learning models in macroeconomics typically focus on either representative agent environments, or environments in which agents are assumed to be unaware of the learning efforts of other agents.⁴ Here we study an explicitly *interactive* learning environment, where agents are aware that other agents are learning. Consequently, part of an agent's Perceived Law of Motion (PLM) must include beliefs about how other agents are learning. We follow the traditional learning literature and assume that while these beliefs can be misspecified during the transition, they are statistically confirmed in the limit. Esponda and Pouzo (2015) refer to this as a 'Berk-Nash Equilibrium'. We show that even though agents observe the same data, their beliefs about stationarity do not 'merge'. Although the policymaker uses the TVP model almost always, the forecaster using the constant parameter model will nonetheless be unable to statistically reject his beliefs about the underlying stationarity of the environment, for essentially the same reason it is statistically difficult to distinguish a random walk from a stationary, but very persistent process.⁵

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. We then allow the policymaker to consider multiple models, and examine the implications of Bayesian Model Averaging. To motivate the subsequent convergence analysis, section 3, reports a variety of simulations. These simulations show that self-confirming parameter drift can generate significant (low frequency) asset price volatility. Section 4 describes the key concepts necessary for the analysis. Section 5 shows that the weight on the TVP model eventually converges to one. We first study the case where individual forecasters ignore the presence of other forecasters. We then show that under self-confirming beliefs about other forecasters beliefs, the same Gresham's Law result emerges. Section 6 discusses important properties of the outcome, including persistent disagreement of beliefs, robustness of the results to alternative interpretations of the model space and alternative interpretations of model averaging. Instead of a single decision maker who averages over his own two models, we consider a decentralized situation, in which two different forecasters forward their price forecast to the

⁴Evans and Honkapohja (2001) provide the definitive summary of this literature.

⁵Kurz (1994) and Acemoglu, Chernozhukov, and Yildiz (2016) contain similar results, for similar reasons. Kurz (1994) appeals to nonstationarity, while Acemoglu et. al. (2016) appeal to model uncertainty and lack of identification.

policymaker, who then forms his own forecast by averaging the two. We show that the main conclusions from the single agent case carry over. Section 7 discusses extensions and potential applications, while the Appendix collects proofs of various technical results.

2. BASELINE MODEL

Our analysis is inspired by the previous work of Evans, Honkapohja, Sargent, and Williams (2013). They study a standard cobweb model, in which a single 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. Our paper extends their analysis in a few different directions. First, instead of a cobweb model, where current price depends on previous expectations, we study an asset pricing model, where current price depends on current expectations of future prices. This is arguably a more empirically relevant case in macroeconomics. Second, and more importantly, we apply the tools of large deviations theory to provide a formal convergence analysis of BMA with endogenous data. Third, and perhaps most importantly, we consider an interactive learning environment, in which beliefs about other agents beliefs are explicitly considered. This is arguably more descriptive of the actual forecasting environment confronting macroeconomic policymakers.

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

$$p_t = \delta z_t + \alpha \mathbb{E}_t p_{t+1} + \sigma \epsilon_t \quad (2.1)$$

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

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

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

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

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

It is well known that Rational Expectations versions of this kind of model cannot explain observed asset price volatility (Shiller (1989)). We explain this volatility by assuming that agents must learn about their environment. Of course, the notion that learning might help to explain asset price volatility is hardly new (see, e.g., Timmermann (1996) for an early and influential example). However, early examples were based on least-squares learning, which exhibited asymptotic convergence to the Rational Expectations Equilibrium. This would be fine if volatility appeared to dissipate over time, but there is no evidence for this. In response, a more recent literature has assumed that agents use so-called *constant gain* learning, which discounts old data. This keeps learning alive. For example, Benhabib and Dave (2014) show

that constant gain learning can generate persistent excess volatility, and can explain why asset prices have fat-tailed distributions even when the distribution of fundamentals is thin-tailed.

Our paper builds on the work of Benhabib and Dave (2014). The key parameter in their analysis is the update gain. Not only do they assume it is bounded away from zero, but they restrict it to be *constant*. Following Sargent and Williams (2005), they note that a constant gain can provide a good approximation to the (steady state) gain of an optimal Kalman filtering algorithm. However, they go on to show that the learning dynamics exhibit recurrent escapes from this steady state. This calls into question whether agents would in fact cling to a constant gain in the presence of such instability. Here we allow the agent to effectively employ a time-varying gain, which is not restricted to be nonzero. We do this by supposing that agents *average* between a constant gain and a decreasing/least-squares gain. Evolution of the model probability weights delivers a state-dependent gain. In some respects, our analysis resembles the gain-switching algorithm of Marcat and Nicolini (2003). However, they require the agent to commit to one or the other, whereas we permit the agent to be a Bayesian, and average between the two. Despite the fact that our specification of the gain is somewhat different, like Benhabib and Dave (2014), we rely on the theory of large deviations to provide an analytical characterization of the escape dynamics.

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

$$p_t = \beta_t z_t + \sigma \epsilon_t \quad (2.4)$$

$$\beta_t = \beta_{t-1} + \sigma_v v_t \quad (2.5)$$

where it is assumed that $\text{cov}(\epsilon, v) = 0$. Note that the Rational Expectations equilibrium is a special case of this, with

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

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

$$\mathcal{M}_0 : \sigma_v^2 = 0.$$

Given this belief, he estimates the unknown parameter of his model using the following Kalman filter algorithm

$$\hat{\beta}_{t+1} = \hat{\beta}_t + \left(\frac{\Sigma_t}{\sigma^2 + \Sigma_t z_t^2} \right) z_t (p_t - \hat{\beta}_t z_t) \quad (2.6)$$

$$\Sigma_{t+1} = \Sigma_t - \frac{(z_t \Sigma_t)^2}{\sigma^2 + \Sigma_t z_t^2} \quad (2.7)$$

where we adopt the common assumption that $\hat{\beta}_t$ is based on time- $(t-1)$ information, while the time- t forecast of p_{t+1} , $\rho \hat{\beta}_t z_t$, can incorporate the latest z_t observation. This assumption is made to avoid simultaneity between beliefs and observations.⁶ The process, Σ_t , represents the agent's evolving estimate of the variance of $\hat{\beta}_t$.

Notice that given his beliefs that parameters are constant, Σ_t converges to zero at rate t^{-1} . This makes sense. If parameters really are constant, then each new observation contributes less and less relative to the existing stock of knowledge. On the other hand, notice that during the transition, the agent's beliefs are inconsistent with the data. He *thinks* β is constant, but due

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

to expectational feedback, his own learning causes β to be time-varying. This can be seen by substituting the agent's time- t forecast into the true model in (2.1)

$$\begin{aligned} p_t &= [\delta + \rho\alpha\hat{\beta}_t]z_t + \sigma\epsilon_t \\ &= T(\hat{\beta}_t)z_t + \sigma\epsilon_t \end{aligned}$$

Opinions differ as to whether this inconsistency is important. As long as the T -mapping between beliefs and outcomes has the appropriate stability properties, the agent's incorrect beliefs will eventually be corrected. Learning-induced parameter variation eventually dissipates, and the agent eventually learns the Rational Expectations equilibrium. However, as pointed out by Bray and Savin (1986), in practice this convergence can be slow, and one might then ask why agents aren't able to detect the parameter variation that their own learning generates. If they do, wouldn't they want to revise their learning algorithm, and if they do, will learning still take place?⁷

This debate is largely academic, however, since the more serious problem with this model is that it fails to explain the data. Since learning is transitory, so is any learning induced parameter instability. Although there is some evidence in favor of a 'Great Moderation' in the volatility of macroeconomic aggregates (at least until the recent financial crisis!), there is little or no evidence for such moderation in asset markets. As a result, more recent work assumes agents view parameter instability as a permanent feature of the environment.

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

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

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)). Still, due to feedback, there is a sense in which his beliefs about parameter instability can be self-confirming, since ongoing belief revisions will produce ongoing parameter instability.

The belief that $\sigma_v^2 > 0$ produces only a minor change in the Kalman filtering algorithm in (2.6) and (2.7). We just need to replace the Riccati equation in (2.7) with the new Riccati equation

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

The additional σ_v^2 term causes Σ_t to now converge to a strictly positive limit, $\bar{\Sigma} > 0$. As noted by Benveniste et. al. (1990, pgs. 139-40), 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 $M_z = E(z_t^2)$ denotes the second moment of the fundamentals process. In addition, if we further assume that

⁷McGough (2003) addresses this issue. He pushes the analysis one step back, and shows that if agents start out with a time-varying parameter learning algorithm, but have priors that this variation damps out over time, then agents can still eventually converge to a constant parameter Rational Expectations equilibrium.

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

$$\hat{\beta}_{t+1} = \hat{\beta}_t + \gamma M_z^{-1} z_t (p_t - \hat{\beta}_t z_t) \quad (2.9)$$

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.

Remember that since the underlying parameters are constant, the agent’s model is ‘wrong’. Wouldn’t a smart agent eventually discover this?⁸ On the one hand, this is an easy question to answer. Since his prior dogmatically rules out the ‘right’ constant parameter model, there is simply no way the agent can ever detect his misspecification, even with an infinite sample. On the other hand, due to the presence of expectational feedback, a more subtle question is whether the agent’s beliefs about parameter instability could become ‘self-confirming’ (Sargent (2008))? 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, chpt. 6), in self-referential models the gain should *not* be treated as a free parameter. It is an equilibrium object. This is because the optimal gain depends on the volatility of the data, but at the same time, the volatility of the data depends on the gain. As in a Rational Expectation Equilibrium, we have a fixed point problem.

In a prescient paper, Evans and Honkapohja (1993) addressed the problem of computing this fixed point. They posed the problem as one of computing a Nash equilibrium. In particular, they ask - Suppose everyone else is using a given gain parameter, so that the data-generating process is consistent with this gain. Would an individual agent have an incentive to switch to a *different* gain? Under appropriate stability conditions, one can then compute the equilibrium gain by iterating on a best response mapping as usual. Later we exploit this idea to study the stability of our more complex BMA 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 this can be written as

$$\hat{\beta}_{t+1} = \hat{\beta}_t + \gamma M_z^{-1} z_t \left\{ [\delta + (\alpha\rho - 1)\hat{\beta}_t] z_t + \sigma\epsilon_t \right\} \quad (2.10)$$

Let $\beta^* = \delta/(1 - \alpha\rho)$ denote the Rational Expectations equilibrium. Also let $\tau_t = t \cdot \gamma$, and then define $\beta(\tau_t) = \hat{\beta}_t$. We can then form the piecewise-constant continuous-time interpolation, $\beta(\tau) = \beta(\tau_t)$ for $\tau \in [t\gamma, t\gamma + \gamma)$. Although for a fixed γ (and σ_v^2) the paths of $\beta(\tau)$ are not continuous, they converge to the following continuous limit as $\sigma_v^2 \rightarrow 0$ (see Evans and Honkapohja (2001) for a proof)

Proposition 2.1. *As $\sigma_v^2 \rightarrow 0$, $\beta(\tau)$ converges weakly to the solution of the following diffusion process*

$$d\beta = -(1 - \alpha\rho)(\beta - \beta^*)d\tau + \gamma M_z^{-1/2} \sigma dW_\tau \quad (2.11)$$

⁸Of course, a constant gain model could be the ‘right’ model too, if the underlying environment features exogenously time-varying elements. After all, it is this possibility that motivates their use in the first place. Interestingly, however, most existing applications of constant gain learning feature environments in which doubts about parameter stability are entirely in the head of the agent.

where dW_τ is the standard Wiener process.

This is an Ornstein-Uhlenbeck process, which generates a stationary Gaussian distribution centered on the Rational Expectations equilibrium, β^* . Notice that the innovation variance is consistent with the agent's priors, since $\gamma^2 \sigma^2 M_z^{-1} = \sigma_v^2$. However, notice also that $d\beta$ is autocorrelated. That is, β does *not* follow a random walk. Strictly speaking then, the agent's priors are misspecified. However, remember that traditional definitions of self-confirming equilibria presume agents have access to infinite samples. In practice, agents only have access to finite samples. Given this, we can ask whether the agent could statistically reject his prior.⁹ This will be difficult when the drift in (2.11) is small. This is the case when: (1) Estimates are close to the β^* , (2) Fundamentals are persistent, so that $\rho \approx 1$, and (3) Feedback is strong, so that $\alpha \approx 1$.

2.4. Model averaging. Dogmatic priors (about anything) are rarely a good idea. So now suppose agents hedge their bets by entertaining the possibility that parameters are constant. Forecasts are then constructed using a traditional Bayesian Model Averaging (BMA) strategy. This strategy effectively ‘convexifies’ the model space.

There are two ways to think about model averaging. One is to assume the competing models are in the mind of a single agent. This is the interpretation in Evans, Honkapohja, Sargent, and Williams (2013). However, if this is the case, it is hard to understand why the agent does not just 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. The second way to think about model averaging is from a more decentralized perspective, where multiple agents construct and revise models, which are then ‘marketed’ to a single 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.¹⁰

We are going to consider both possibilities, mainly for pedagogical reasons. This is because the first approach is easier to formalize, since it just involves specifying the beliefs of a single agent. In contrast, with multiple agents and multiple models, one must specify how agents perceive the forecasting efforts of other agents. If agents are aware that forecasts are being used by a policymaker, whose actions potentially influence the data-generating process, they must then form beliefs over *other* forecasters’ beliefs. For example, even if you are convinced the underlying environment is stationary, if you know a policymaker is putting some weight on a TVP model, then your own model must allow for the induced nonstationarity. Of course, specifying beliefs about others’ beliefs is fraught with difficulties. In Bayesian settings it quickly produces an infinite regress, which can only be tamed by introducing strong coordination assumptions, like common priors. Rather than pursue this route, we introduce an element of bounded rationality, by allowing agents’ beliefs about other agents’ beliefs to be misspecified. However, as in Esponda and Pouzo (2015), we impose discipline by requiring beliefs to be statistically confirmed in the limit. As it turns out, the asymptotic properties of the model averaging dynamics are the same for both the single agent and multiple agent environments. We start with the former, as it is easier. Once the argument is clear for this case, the more complex interactive learning environment will be easier to understand.

Let π_t denote the current probability assigned by the policymaker to \mathcal{M}_1 , the TVP model, and let $\beta_t(i)$ denote the current parameter estimate for \mathcal{M}_i . The policymaker’s time- t forecast

⁹In the language of Hansen and Sargent (2008), we can compute the detection error probability.

¹⁰Timmermann (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.

becomes¹¹

$$\mathbf{E}_t p_{t+1} = \rho[\pi_t \beta_t(1) + (1 - \pi_t) \beta_t(0)] z_t \quad (2.12)$$

Substituting this into the actual law of motion for prices implies that parameter estimates evolve according to

$$\beta_{t+1}(i) = \beta_t(i) + \left(\frac{\Sigma_t(i)}{\sigma^2 + \Sigma_t(i) z_t^2} \right) z_t \{ [\delta + \alpha \rho [\pi_t \beta_t(1) + (1 - \pi_t) \beta_t(0)] - \beta_t(i)] z_t + \sigma \epsilon_t \} \quad (2.13)$$

As usual, we suppose the policymaker neglects the feedback from his forecast to the actual price process. Note that the only difference between the two parameter update equations arises from their gain sequences, $\Sigma_t(i)$. With a single forecaster who neglects feedback, these two gain sequences are independent of model averaging.

3. SIMULATIONS

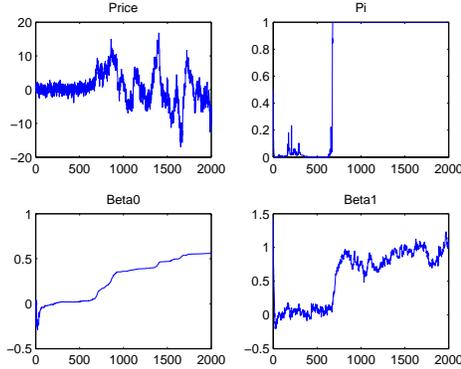
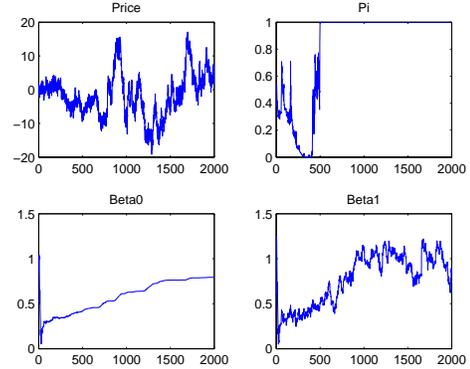
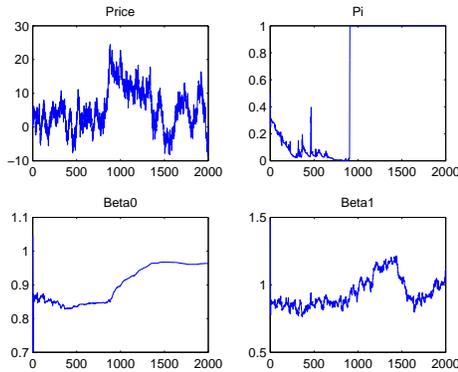
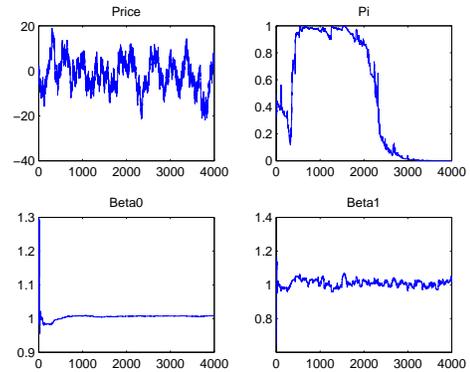
As noted in Section 2, the present-value asset pricing model in (2.1)-(2.2) has been subjected to a lot of previous empirical work, mostly with negative results. Perhaps its biggest problem is its failure to generate sufficient volatility (Shiller (1989)). Our results suggest that this negative assessment could be premature. To examine this possibility, we calibrate the model using parameter values that have been used in the past, and see whether this can generate the sort of self-confirming volatility that our analysis suggests is possible.

Most of the parameters are easy to calibrate. We know observed fundamentals are persistent, so we set $\rho = .99$. Remember, the agent is assumed to know this. Similarly, we know discount factors are close to 1, so we set the feedback parameter to $\alpha = .96$. Since δ depends on units, we just normalize it by setting $\delta = (1 - \alpha\rho)$. This implies the self-confirming 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 ($\sigma^2 = \sigma_z^2 = 1$). That leaves one remaining free parameter, σ_v^2 . Of course, this is a crucial parameter, since it determines the agent's prior beliefs about parameter instability. If it is too big, then the TVP model will be at a big disadvantage during tranquil times, and will therefore have a difficult time displacing the constant parameter model. On the other hand, if it's too small, self-confirming volatility will be empirically irrelevant.

Figures 1-4 report typical simulations for three alternative values, $\sigma_v^2 = (.0005, .0001, .00001)$. Figures 1 and 2 are for the value $\sigma_v^2 = .0005$. 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 $\gamma = .07$, which is quite typical of values used in prior empirical work. These figures also illustrate a typical feature of the sample paths when σ_v^2 is relatively high, i.e., convergence to one or the other boundaries occurs relatively quickly, usually by around $T = 500$.

Figures 3-4 use smaller values of σ_v^2 . Generally speaking, smaller values of σ_v^2 delay convergence. In Figure 3, where $\sigma_v^2 = .0001$, convergence to $\pi = 1$ once again takes place, but now its price volatility implications are not quite so dramatic. Volatility is only 41.7% higher when $\pi = 1$. Once again, the implied steady-state gain ($\gamma = .03$) is typical of values used in empirical work. Figure 4 uses a still smaller prior variance, $\sigma_v^2 = .00001$. Now the two models do not differ by much. Steady-state price volatility is only 13% higher when $\pi = 1$. Notice that

¹¹To ease notation in what follows, we shall henceforth omit the hats from the parameter estimates.

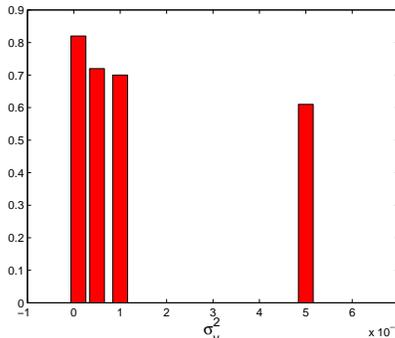
Figure 1: $\sigma_v^2 = .0005$ Figure 2: $\sigma_v^2 = .0005$ Figure 3: $\sigma_v^2 = .0001$ Figure 4: $\sigma_v^2 = .00001$

because the two models are so similar, it becomes easier to escape the $\pi = 1$ equilibrium. Since the TVP world is not that volatile, a constant parameter model does not do that badly.¹²

The one feature that is perhaps not accurately portrayed by these figures is the fact that on empirically relevant time-scales convergence to either boundary can occur. This fact was emphasized by Evans, Honkapohja, Sargent, and Williams (2013). Although our previous results imply that eventually the $\pi = 1$ equilibrium will dominate, our simulations indicate that the $\pi = 0$ equilibrium can persist for a *long* time. For example, we conducted 10,000 simulations, each of length $T = 2000$, and counted the proportion of times convergence to $\pi = 1$ occurred for various values of σ_v^2 . As above, the simulations were initialized at $\pi = 0.5$, with small random perturbations of the coefficients around their self-confirming equilibrium values. Figure 5 displays the results,

The probability of convergence to $\pi = 1$ declines with σ_v^2 . For $\sigma_v^2 = .1 \times 10^{-4}$, convergence occurs more than 80% of the time, whereas for the benchmark value used above, $\sigma_v^2 = 5 \times 10^{-4}$,

¹²Notice that Figure 4 uses a $T = 4000$ sample length, while Figures 1-4 use $T = 2000$. As σ_v^2 decreases, things evolve more slowly, so it becomes necessary to expand the simulation length.

Figure 5: Probability of convergence to $\pi = 1$ after 2000 periods

convergence to $\pi = 1$ occurs only about 60% of the time.¹³ As our above analysis makes clear, however, one must exercise some caution when interpreting results like these. In Figure 5, ‘convergence’ was simply defined as the value of π at the end of each simulation run (i.e., at $T = 2000$). According to our results, however, *eventually* π will escape from 1. What our theory actually predicts is that as the sample length becomes infinitely long, the *proportion of time* spent near $\pi = 1$ goes to unity. It doesn’t imply that π never returns to 0. Although escapes from $\pi = 1$ are more likely to occur relatively early in the game, before $\beta_t(0)$ has settled down, as σ_v^2 decreases, escapes can occur relatively late as well. Figure 4 nicely illustrates this possibility. Hence, the results in Figure 5 are merely meant to convey the possibility that on realistic time-scales, ‘convergence’ to either $\pi = 1$ or to $\pi = 0$ can occur.

The fact that convergence to either $\pi = 0$ or $\pi = 1$ can occur on relevant time-scales is interesting, since it suggests that whether we live in a tranquil or volatile economy is somewhat random and history-dependent. It also highlights a potentially adverse long-term effect of ‘large deviation’ events, like financial crises. Although being alert to the possibility of financial crises is probably a good thing on net, if it makes individuals living in a less than fully understood self-referential environment more reactive, it could create its own problems.

4. PRELIMINARIES

Since our analysis relies heavily on time-scale separation methods, we begin with a brief discussion of time scales.

4.1. Time scales. Since each component of $(\beta_t(1), \beta_t(0), \pi_t)$ evolves at a different “speed,” it is important to define clearly a notion of speed in terms of a benchmark time scale. We use the sample average time scale, t^{-1} , as the benchmark. This is the time-scale at which $\beta_t(0)$ evolves. More precisely, $\forall \tau > 0$, we can find the unique integer satisfying

$$\sum_{k=1}^{K-1} \frac{1}{k} < \tau < \sum_{k=1}^K \frac{1}{k}.$$

Let $m(\tau) = K$ and define

$$t_K = \sum_{k=1}^K \frac{1}{k}$$

¹³For $\sigma_v^2 = 0$ convergence should occur 50% of the time.

Therefore, $t_K \rightarrow \infty$ as $K \rightarrow \infty$. We are interested in sample paths over the tail interval $[t_K, t_K + \tau)$. Hence, we are interested in the speed of evolution in the right hand tail of a stochastic process.

It is important to recall that the function m which maps “fictitious time” to a number of periods is defined for the sample average time scale, and that we are interested in the right tail of the sequence.

Definition 4.1. *Let φ_t be a stochastic process. We say that φ_t evolves on a faster time scale than $\beta_t(0)$ if $\forall \tau > 0$,*

$$\lim_{K \rightarrow \infty} \sum_{t=K}^{m(t_K + \tau)} |\varphi_t - \varphi_{t-1}| = \infty,$$

with a positive probability, and evolves on a slower time scale than $\beta_t(0)$ if

$$\lim_{K \rightarrow \infty} \sum_{t=K}^{m(t_K + \tau)} |\varphi_t - \varphi_{t-1}| = 0$$

with probability 1. We say that φ_t evolves on the same time scale as $\beta_t(0)$ if φ_t does not evolve at a faster or slower time scale.

If a stochastic process φ_t has a recursive representation

$$\varphi_t = \varphi_{t-1} + \varepsilon_t Y_t(\varphi_{t-1}, X_t)$$

where X_t is the state variable, Y_t is a uniformly bounded function, and $\varepsilon_t > 0$ is the gain function, we can compare time scales by comparing gain functions. Consider another recursive formula

$$\tilde{\varphi}_t = \tilde{\varphi}_{t-1} + \tilde{\varepsilon}_t \tilde{Y}_t(\tilde{\varphi}_{t-1}, X_t)$$

with uniformly bounded \tilde{Y}_t and gain function $\tilde{\varepsilon}_t > 0$. One can show that φ_t evolves on a faster time scale than $\tilde{\varphi}_t$ if

$$\lim_{t \rightarrow \infty} \frac{\tilde{\varepsilon}_t}{\varepsilon_t} = 0.$$

Hence, if a stochastic process can be represented recursively, it makes sense to represent its time scale or speed of evolution in terms of its gain sequence.

4.2. Odds ratio. The policymaker updates his beliefs about models using Bayes Rule. After a long tedious calculation, the Bayesian updating scheme for π_t can be written as (see Evans, Honkapohja, Sargent, and Williams (2013) for a partial derivation)

$$\frac{1}{\pi_{t+1}} - 1 = \frac{A_{t+1}(0)}{A_{t+1}(1)} \left(\frac{1}{\pi_t} - 1 \right) \quad (4.14)$$

where

$$A_t(i) = \frac{1}{\sqrt{2\pi \text{MSE}(i)}} e^{-\frac{(p_t - \hat{p}_t(i))^2}{2\text{MSE}(i)}} \quad (4.15)$$

is the time- t predictive likelihood function for model \mathcal{M}_i , and

$$\text{MSE}(i) = \text{E}(p_t - \hat{p}_t(i))^2$$

is the mean-squared forecast error of \mathcal{M}_i . Remember that $\hat{p}_t(i)$ depends on assumptions about the nature of model averaging. In the single-agent case, the multiple models are in the mind of the policymaker, and $\hat{p}_t(i) = \rho \beta_t(i) z_t$. In the multiple-agent case, the policymaker uses the

forecasts provided by other agents, and the $\hat{p}_t(i)$ are given by the more complicated expressions in equations (6.38).

To study the dynamics of π_t it is useful to rewrite eq. (4.14) as follows

$$\pi_{t+1} = \pi_t + \pi_t(1 - \pi_t) \left[\frac{A_{t+1}(1)/A_{t+1}(0) - 1}{1 + \pi_t(A_{t+1}(1)/A_{t+1}(0) - 1)} \right] \quad (4.16)$$

which has the familiar form of a discrete-time replicator equation, with a stochastic, state-dependent, fitness function determined by the likelihood ratio. Equation (4.16) 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 $\mathbb{E}(A_{t+1}(1)/A_{t+1}(0)) = 1$. Later we shall see that this occurs when $\pi = \frac{1}{2\rho\alpha}$, which is interior if feedback is strong enough (i.e., if $\alpha > \frac{1}{2\rho}$). 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. This will become apparent when we turn to the simulations in Section 4. One remaining issue is whether π_t could ever become absorbed at one of the boundary points.

Proposition 4.2. *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. See Appendix A. □

Since the distributions here are assumed to be Gaussian, they obviously have full support, so Proposition 4.2 applies. Although the boundary points are unattainable in finite time, the replicator equation for π_t in eq. (4.16) makes it clear that π_t will spend most of its time near these boundary points, since the relationship between π_t and π_{t+1} has the familiar logit function shape, which flattens out near the boundaries. As a result, π_t evolves very slowly near the boundary points. In fact, we shall now show that it evolves even more slowly than the t^{-1} time-scale of $\beta_t(0)$. This means that when studying the dynamics of the coefficient estimates near the boundaries, we can treat π_t as fixed.

4.3. 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)

$$\dot{\phi} = \mathbb{E} \left[\ln \frac{A_t(0)}{A_t(1)} \right] - \phi$$

which has a unique stable point

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

Note that if $\phi^* > 0$, $\pi_t \rightarrow 0$, while if $\phi^* < 0$, $\pi_t \rightarrow 1$. The focus of the ensuing analysis is to identify the conditions under which π_t converges to 1, or 0. Thus, the sign of ϕ^* , rather than its value, is an important object of investigation.

We show that π_t evolves even more slowly than the t^{-1} time-scale of $\beta_t(0)$ and ϕ_t . This means that when studying the dynamics of the coefficient estimates near the boundaries, we can treat π_t as fixed. In order to make the notion of “more slowly” precise, we need to define precisely the time scale.

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

Lemma 4.3.

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

and π_t evolves at a slower time scale than $1/t$.

Proof. See Appendix B. □

Therefore, without loss of generality, we can assume that $\pi = \{0, 1\}$ are the only limit points of $\{\pi_t\}$. After renumbering a convergent subsequence, suppose $\pi_t \rightarrow 1$. Following the same reasoning as in the proof of Lemma 4.3, we can prove that

$$\beta_t(0) \rightarrow \frac{\delta}{1 - \alpha\rho}$$

with probability 1, and

$$\beta_t(1) \rightarrow \frac{\delta}{1 - \alpha\rho}$$

weakly.

If $\alpha\rho > 1/2$, and π_t close to 1, then a simple calculation shows that

$$\mathbb{E} \ln \frac{A_{t+1}(0)}{A_{t+1}(1)} < 0$$

which implies that

$$\phi_t \rightarrow \mathbb{E} \ln \frac{A_{t+1}(0)}{A_{t+1}(1)} < 0.$$

with probability 1. Given $\beta_t(0) = \beta_t(1) = \delta/(1 - \alpha\rho)$, $\pi_t \rightarrow 1$ with probability 1, proving that $\pi_t = 1$ is locally attracting. Similarly, we can show that if π_t is close to 0, then

$$\mathbb{E} \ln \frac{A_t(0)}{A_t(1)} > 0.$$

Following the same argument, we prove that $\pi_t = 0$ is locally attracting.

In general, the speed of evolution of π_t compared to $\beta_t(i)$ ($i = 1, 2$) is difficult to compute. But, in the neighborhood of the boundaries, we can show that π_t evolves on an even slower time scale than $\beta_t(0)$.

Lemma 4.4.

$$\mathbb{P} \left(\exists \{\phi_t\}, \lim_{k \rightarrow \infty} \phi_t = 0 \right) = 0$$

and π_t evolves at a slower time scale than $1/t$ (or $\beta_t(0)$) with probability 1.

Proof. See Appendix C □

Lemma 4.4 asserts that in the neighborhood of the boundaries, the hierarchy of time scales is such that $\beta_t(1)$ evolves at a faster time scale than $\beta_t(0)$, while $\beta_t(0)$, which evolves on the same time scale as ϕ_t , evolves at a faster time scale than π_t . This time scale hierarchy greatly facilitates the analysis of escape dynamics.

5. MODEL AVERAGING DYNAMICS

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

$$\beta_{t+1}(0) = \beta_t(0) + \left(\frac{\Sigma_t(0)}{\sigma^2 + \Sigma_t(0)z_t^2} \right) z_t(p_t - \beta_t(0)z_t) \quad (5.17)$$

$$\Sigma_{t+1}(0) = \Sigma_t(0) - \frac{(z_t \Sigma_t(0))^2}{\sigma^2 + \Sigma_t(0)z_t^2} \quad (5.18)$$

$$\beta_{t+1}(1) = \beta_t(1) + \left(\frac{\Sigma_t(1)}{\sigma^2 + \Sigma_t(1)z_t^2} \right) z_t(p_t - \beta_t(1)z_t) \quad (5.19)$$

$$\Sigma_{t+1}(1) = \Sigma_t(1) - \frac{(z_t \Sigma_t(1))^2}{\sigma^2 + \Sigma_t(1)z_t^2} + \sigma_v^2 \quad (5.20)$$

$$\frac{1}{\pi_{t+1}} - 1 = \frac{A_{t+1}(0)}{A_{t+1}(1)} \left(\frac{1}{\pi_t} - 1 \right) \quad (5.21)$$

where $A_t(i)$ is defined as (4.15) $\forall i, t$. The actual law of motion for p_t is

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

Except for (5.21), $(\beta_t(0), \Sigma_t(0), \beta_t(1), \Sigma_t(1))$ evolve according to stochastic difference equations. A standard method to investigate the asymptotic properties of a stochastic difference equation is to first calculate the dynamics of the mean, and to then infer the asymptotic properties of the stochastic variables from the mean dynamics (Kushner and Yin (1997)).

However, this “standard” approach does not apply here, because $\beta_t(1), \beta_t(0)$ and π_t evolve at different speeds. One can easily show that $\Sigma_t(0) \rightarrow 0$, while $\Sigma_t(1) \rightarrow \bar{\Sigma} > 0$. For large values of t , $\beta_t(1)$ evolves infinitely faster than $\beta_t(0)$, since the gain function of $\beta_t(1)$ is uniformly bounded away from 0, while that of $\beta_t(0)$ vanishes. Lemma 4.4 shows that π_t evolves more slowly than $\beta_t(0)$. To accommodate the different “speeds” of evolution of these variables,

we use multiple time-scale stochastic approximation methods (Borkar (2008)). Exploiting the time-scale hierarchy, we can first characterize the asymptotic properties of $\beta_t(1)$, for fixed values of $(\beta_t(0), \pi_t)$. After calculating the invariant distribution of $\beta_t(1)$, we then investigate the asymptotic properties of $\beta_t(0)$, and then finally those of π_t .

5.1. Convergence. Since the proof is very much in line with the second part of the proof of Lemma 4.3, we just state the result without proof.

Proposition 5.1. $\lim_{t \rightarrow \infty} \beta_t(1) = \frac{\delta}{1-\alpha\rho}$ weakly, while $\lim_{t \rightarrow \infty} \beta_t(0) = \frac{\delta}{1-\alpha\rho}$ with probability 1. $\pi_t \rightarrow \{0, 1\}$ with probability 1, as $t \rightarrow \infty$.

It is helpful to calculate the domain of attraction of each locally stable point, from the perspective of $\beta_t(0)$, assuming that $\beta_t(1)$ has already reached its own stationary distribution. Comparing the likelihoods, we can compute the domain of attraction for $(\pi, \beta(0), \beta(1)) = (0, \frac{\delta}{1-\rho\alpha}, \frac{\delta}{1-\alpha\rho})$ is the interior of

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

which is roughly the area of

$$\left\{ (\pi, \beta(0), \beta(1)) \mid \left(\beta(0) - \frac{\delta}{1-\alpha\rho} \right)^2 < (1 - 2\alpha\rho\pi)\sigma_\xi^2 \left(\frac{1-\alpha\rho\pi}{1-\alpha\rho} \right)^2 \right\}$$

where the mean-squared forecast error of \mathcal{M}_0 is smaller than that of \mathcal{M}_1 . The difference arises from the fact that the expected likelihood ratio differs from the ratio of expected mean-squared forecast errors.

The interior of the complement of \mathcal{D}_0 is the domain of attraction for $(\pi, \beta(0), \beta(1)) = (1, \frac{\delta}{1-\rho\alpha}, \frac{\delta}{1-\alpha\rho})$

For small $\sigma_v > 0$, one can imagine \mathcal{D}_0 as a narrow ‘‘cone’’ in the space of $(\beta(0), \pi)$, with its apex at $(\beta(0), \pi) = \left(\frac{\delta}{1-\alpha\rho}, \frac{1}{2\alpha\rho} \right)$ and its base along the line $\pi = 0$, where $\beta(0)$ is in $\left[\frac{\delta}{1-\alpha\rho} - \frac{\sigma_\xi}{1-\alpha\rho}, \frac{\delta}{1-\alpha\rho} + \frac{\sigma_\xi}{1-\alpha\rho} \right]$. Figure 6 plots \mathcal{D}_0 for the baseline parameter values used in the previous simulations. The formal analysis will make the notion of being ‘‘narrow’’ precise.

5.2. Escape probabilities. We have three endogenous variables $(\pi_t, \beta_t(0), \beta_t(1))$, which converge to one of the two locally stable points: $(0, \delta/(1-\alpha\rho), \delta/(1-\alpha\rho))$ or $(1, \delta/(1-\alpha\rho), \delta/(1-\alpha\rho))$. Let us identify a specific stable point by the value of π_t at the stable point. Similarly, let \mathcal{D}_0 be the domain of attraction to $\pi_t = 0$, and \mathcal{D}_1 be the domain of attraction to $\pi_t = 1$.

For fixed $\sigma_v^2 > 0$, the distribution of $(\pi_t, \beta_t(0), \beta_t(1))$ assigns a large weight to either of the two locally stable points as $t \rightarrow \infty$. Our main interest is the limit of this probability distribution as $\sigma_v^2 \rightarrow 0$.

To calculate the limit probability distribution, we need to calculate the probabilities that $(\pi_t, \beta_t(0), \beta_t(1))$ escape from each domain of attraction. Standard results from large deviation theory imply that the escape probabilities can be parameterized by their large deviation rate functions.

Lemma 5.2. *There exists $r_0 \in [0, \infty]$ so that*

$$- \lim_{\sigma_v^2 \rightarrow 0} \lim_{t \rightarrow \infty} \frac{1}{t} \log \mathbf{P} \left(\exists t, (\pi_t, \beta_t(0), \beta_t(1)) \in \mathcal{D}_1 \mid (\pi_1, \beta_1(0), \beta_1(1)) = \left(0, \frac{\delta}{1-\alpha\rho}, \frac{\delta}{1-\alpha\rho} \right) \right) = r_0$$

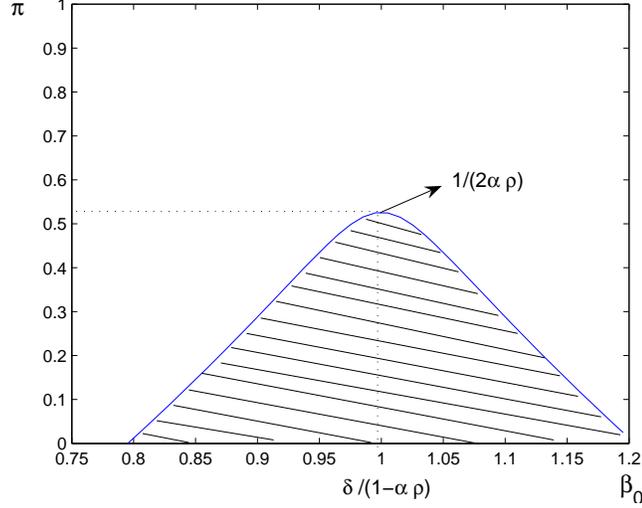


Figure 6: \mathcal{D}_0 , the domain of Attraction for $\pi = 0$. In simulations presented in section 3, we normalize the parameters so that $\delta/(1 - \alpha\rho) = 1$. The shaded region is “narrow” if we use the same scale for the horizontal axis as the scale of the vertical axis.

and $\exists r_1 \in [0, \infty]$ so that

$$- \lim_{\sigma_v^2 \rightarrow 0} \lim_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P} \left(\exists t, (\pi_t, \beta_t(0), \beta_t(1)) \in \mathcal{D}_0 \mid (\pi_1, \beta_1(0), \beta_1(1)) = \left(1, \frac{\delta}{1 - \alpha\rho}, \frac{\delta}{1 - \alpha\rho} \right) \right) = r_1.$$

The large deviation parameter r_i ($i = 1, 2$) quantifies how difficult it is to escape from \mathcal{D}_i , with $r_i = \infty$ meaning that the escape never occurs, and $r_i = 0$ meaning that the escape occurs immediately with probability 1 (on a logarithmic time-scale).

To calculate the relative duration times of $(\pi_t, \beta_t(0), \beta_t(1))$ around each locally attractive boundary point, we need to compute the following ratio

$$\lim_{\sigma_v^2 \rightarrow 0} \lim_{t \rightarrow \infty} \frac{\mathbb{P} \left(\exists t, (\pi_t, \beta_t(0), \beta_t(1)) \in \mathcal{D}_0 \mid (\pi_0, \beta_0(0), \beta_0(1)) = \left(1, \frac{\delta}{1 - \alpha\rho}, \frac{\delta}{1 - \alpha\rho} \right) \right)}{\mathbb{P} \left(\exists t, (\pi_t, \beta_t(0), \beta_t(1)) \in \mathcal{D}_1 \mid (\pi_0, \beta_0(0), \beta_0(1)) = \left(0, \frac{\delta}{1 - \alpha\rho}, \frac{\delta}{1 - \alpha\rho} \right) \right)}.$$

Note that $(\pi_t, \beta_t(0), \beta_t(1))$ stays in the neighborhood of $\left(1, \frac{\delta}{1 - \alpha\rho}, \frac{\delta}{1 - \alpha\rho} \right)$ almost always in the limit, if $r_0 < r_1$, and vice versa.

Proposition 5.3. $\exists \sigma'_v > 0$ such that $\forall \sigma_v \in (0, \sigma'_v)$, $r_1 > r_0$.

Proof. See Appendix D. □

For fixed $\sigma_v^2 > 0$, the asymptotic distribution of π_t assigns a large weight to 1 and 0, since 1 and 0 are locally stable points. Between the two locally stable points, we are interested in which locally stable point is more salient than the other. One way to determine which equilibrium is more salient than the other would be to compare the amount of time when π_t stays in a small neighborhood of each locally stable point. For fixed T , σ_v^2 and $\varepsilon > 0$, define

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

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

As a corollary of Proposition 5.3, we can show that for a small $\sigma_v^2 > 0$, π_t stays in the neighborhood of 1 almost always.

Theorem 5.4. $\forall \varepsilon > 0$,

$$\lim_{\sigma_v^2 \rightarrow 0} \lim_{T \rightarrow \infty} \mathbf{E} \frac{T_1}{T} = 1$$

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

6. DISCUSSION

6.1. Persistent disagreement. We have assumed that the only agent making any real decisions here is the policymaker. He must decide how much weight to assign to each model. What about the forecasters? We have implicitly assumed they dogmatically stick to their own beliefs about stationarity. However, since they both observe the same data, one might expect from standard merging arguments that their opposing views would ultimately disappear, as long as their priors put at least some weight on the alternative. Strictly speaking, however, no finite amount of data could ever convince either that their beliefs about stationarity were incorrect (Kurz (1994)). Of course, complete certainty is perhaps not a reasonable standard. Could either be statistically convinced to change his views about stationarity?

Given that π_t ultimately converges to unity, both forecasters end up with misspecified models, but 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 REE value. \mathcal{M}_0 thinks parameters are constant, but when $\pi = 1$, they instead fluctuate randomly around the REE value. Hence, the \mathcal{M}_1 forecaster confronts the problem of rejecting a unit root null hypothesis. When $\alpha \approx 1$ and fundamentals are persistent, the alternative is close to being nonstationary. As is well known, the power of unit root tests is notoriously low. It can take thousands of observations before rejection occurs with any confidence. On the other hand, the \mathcal{M}_0 forecaster 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, especially if $\sigma_v > 0$ is small.

Hence, the two forecasters ‘agree to disagree’ about stationarity for a long time. Although this disagreement would disappear with an infinite sample, as $\sigma_v \rightarrow 0$, it takes 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. However, even a ‘grain of doubt’ opens the door to the ultimate dominance of \mathcal{M}_1 .

6.2. Links to 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 is Bansal and Yaron's (2004) long-run risk model. They note that the combination of Epstein-Zin preferences and a small-noise 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 random walk 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.

6.3. Averaging vs. Selection. Theorem 5.4 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. Of course, 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.

An alternative strategy would be to select a particular model using a well defined selection criterion. For example, if the policymaker were simply to flip a coin, and commit to a particular model, there would be a 50/50 chance that \mathcal{M}_0 would be used, which is certainly better odds than with model averaging, where \mathcal{M}_1 is used with probability 1! Of course, committing to a model on the basis of a coin flip is a rather questionable strategy, to say the least. At a minimum, any reasonable strategy should be sequential, and allow the policymaker the freedom to reconsider his choice based on observed outcomes.

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

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

$$\beta_t(i) = \beta_{t-1}(i) + \eta_t(i)\Lambda_t(i) \tag{6.22}$$

$$\Lambda_t(i) = z_{t-1}[p_t - z_t\beta_{t-1}(i)] \tag{6.23}$$

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

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

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

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

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

$$\theta_t(i) = \theta_{t-1}(i) + \eta_t^\alpha(i) \left[\frac{\Lambda_t^2}{\Omega_t(i)} - \theta_{t-1}^i \right] \quad (6.24)$$

$$\Omega_t(i) = \Omega_{t-1}(i) + \eta_t^\alpha(i) [\Lambda_t^2(i) - \Omega_{t-1}(i)] \quad (6.25)$$

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 $\bar{\theta}$ 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 $\theta_t(s_{t-1}) < \bar{\theta}$. Otherwise, $s_t = 1$ with probability one half, and $s_t = 0$ with probability one half.¹⁴

Cho and Kasa (2015) demonstrated that in the long run, the policy maker chooses a model with the largest large deviation parameter at the self-confirming equilibrium. In our case, $\beta_t(0)$ and $\beta_t(1)$ have the same large deviation parameter. The large deviation properties of $\beta_t(i)$ are determined by the statistical properties of the forecasting error $\Lambda_t(i)$. When $\sigma_v \rightarrow 0$, the statistical properties of $\Lambda_t(0)$ and $\Lambda_t(1)$ converge, in the sense that the logarithmic moment generating functions of $\Lambda_t(1)$ and $\Lambda_t(0)$ converge. Standard results from the large deviation literature then show that the large deviation parameters of the two models converge (e.g., Dupuis and Kushner (1987)).

Proposition 6.1. *Define*

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

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

$$\lim_{\sigma_v \rightarrow 0} \mathbb{E} \frac{T_1^v}{T} = \frac{1}{2}.$$

Proof. See Cho and Kasa (2015). □

One has to take Proposition 6.1 with a grain of salt, because the convergence speed of the large deviation parameters is known to be very slow. Even though \mathcal{M}_0 is used one half of the time in the limit as $\sigma_v \rightarrow 0$, the frequency that \mathcal{M}_0 is used for a small positive value of $\sigma_v > 0$ can be close to 1.

Here we chose $\bar{\theta} = 2.5$, so that it is very unlikely that single shock realizations produce model rejections, and $\alpha = 1/2$ to speed up the validation process. The numerical results are robust to the selection of these parameters. Figure 7 shows the empirical distribution of s_t for different values of σ_v . We run 1,000 period simulations as many as 1,000 times, and plot the empirical distribution of s_t for $\sigma_v \in \Omega$ where $\Omega = \{.0005, .0004, .0003, .0002, .0001, .00005, .000025, .00001, .000005, .000001\}$. $\sigma_v = 0.0005$ is hardly a large value, but the probability weight assigned by the empirical distribution of s_t at 0 is close to 1 under $\sigma_v = 0.0005$. As $\sigma_v \rightarrow 0$, the distribution of s_t converges to a bimodal distribution concentrated at either 0 or 1, piling a larger weight to $s_t = 1$. Figure 8 shows the empirical expected value of T_1^v/T_v . If $\sigma_v = 0.0005$, T_1^v/T is close to 0. But, as $\sigma_v \rightarrow 0$, the empirical expected value of T_1^v/T converges to 0.5, as the theory predicts.

¹⁴The random selection rule here does not affect the long run distribution of s_t , as shown in Cho and Kasa (2015).

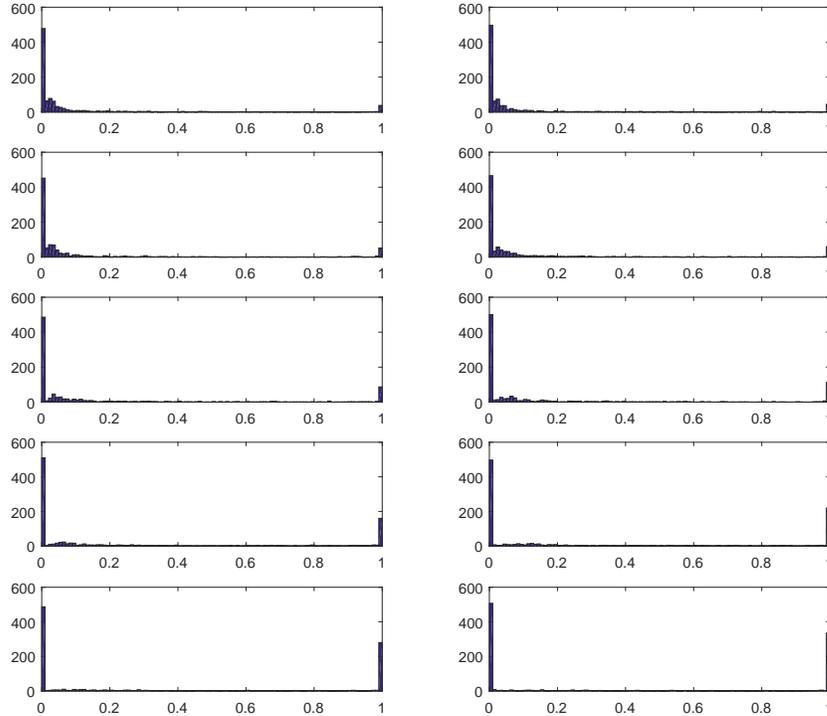


Figure 7: $s_t = 1$ if \mathcal{M}_1 is used by the policy maker in period t , and otherwise, $s_t = 0$. We report empirical distributions of s_t for $\sigma_v \in \Omega$. The number in the vertical axis is 1000 times of the probability. The horizontal axis shows T_1^v/T , which ranges from 0 to 1. The panel is arranged from left to right and from top to bottom, according to the descending order of the values of σ_v . Note that the distribution converges to a bimodal distribution, and that the probability mass at $s_t = 1$ converges to 0.5.

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

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

In this case, there is no model averaging. There is just *one* model, with σ_v^2 viewed as an unknown parameter to be estimated. To address the stability question we exploit the connection discussed in section 2.3 between σ_v^2 and the steady-state gain, γ . Because the data are endogenous, we must employ the macroeconomist's 'big K , little k ' trick, which in our case we refer to as 'big Γ , little γ '. That is, our stability question can be posed as follows: Given that

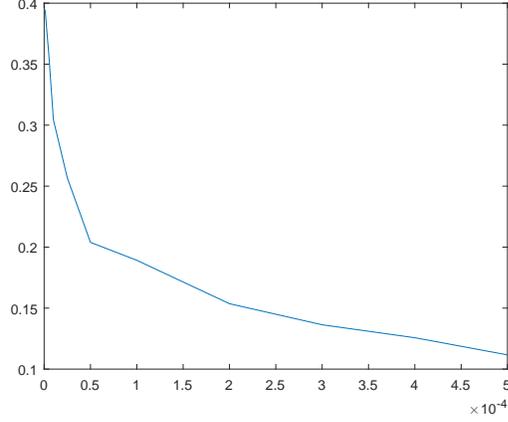


Figure 8: The horizontal axis is the values of σ_v , and the vertical axis is the proportion of time \mathcal{M}_1 is selected. The proportion of time is close to 0, if $\sigma_v = 0.0005$, but converges to 0.5 as $\sigma_v \rightarrow 0$.

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

$$p_t = \delta + \alpha E_t p_{t+1} + \sigma \epsilon_t \quad (6.26)$$

and the agent's perceived model becomes

$$p_t = \beta_t + \sigma \epsilon_t \quad (6.27)$$

$$\beta_t = \beta_{t-1} + \sigma_v v_t \quad (6.28)$$

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. (6.27) and (6.28) represent an example of 'random walk plus noise' model of Muth (1960), in which constant gain updating is optimal. To see this, write p_t as the following ARMA(1,1) process

$$p_t = p_{t-1} + \varepsilon_t - (1 - \Gamma)\varepsilon_{t-1} \quad \Gamma = \frac{\sqrt{4s + s^2} - s}{2} \quad \sigma_\varepsilon^2 = \frac{\sigma^2}{1 - \Gamma} \quad (6.29)$$

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

$$\hat{p}_{t+1} = \hat{p}_t + \Gamma(p_t - \hat{p}_t) \quad (6.30)$$

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

$$\hat{p}_{t+1} = \left(\frac{\Gamma}{1 - (1 - \Gamma)L} \right) p_t \quad (6.31)$$

Substituting this into the true model in eq. (6.26) yields the actual price process as a function of aggregate beliefs

$$\begin{aligned} p_t &= \frac{\delta}{1-\alpha} + \left(\frac{1-(1-\Gamma)L}{1-\left(\frac{1-\Gamma}{1-\alpha\Gamma}\right)L} \right) \frac{\epsilon_t}{1-\alpha\Gamma} \\ &\equiv \bar{p} + f(L;\Gamma)\tilde{\epsilon}_t \end{aligned} \quad (6.32)$$

Now for the ‘big Γ , little γ ’ trick. Suppose prices evolve according eq. (6.32), and that an individual agent has the perceived model

$$\begin{aligned} p_t &= \frac{1-(1-\gamma)L}{1-L} u_t \\ &\equiv h(L;\gamma)u_t \end{aligned} \quad (6.33)$$

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. (6.32). 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 (6.32) and (6.33)¹⁵

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

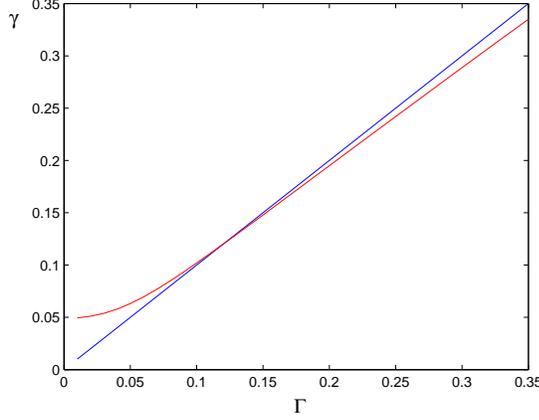
where $F(\omega) = f(e^{-i\omega})f(e^{i\omega})$ and $H(\omega) = h(e^{-i\omega})h(e^{i\omega})$ are the spectral densities of $f(L)$ in eq. (6.32) and $h(L)$ in eq. (6.33). Although this problem cannot be solved with pencil and paper, it is easily solved numerically. Figure 9 plots the best response map using the same benchmark parameter values as before (except, of course, $\rho = 1$ now)¹⁶

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 9 suggests that the best response map intersects the 45 degree line somewhere in the interval (.10, .15). This suggests that the value of σ_v^2 used for the benchmark TVP model in section 4 was a little too small, since it implied a steady-state gain of .072.

6.5. Multiple Agent Model Averaging. Now consider a decentralized setting with three interacting (types of) agents: Two masses of competing, infinitesimal forecasters (\mathcal{M}_0 and \mathcal{M}_1), and a policymaker who combines forecasts. The policymaker is the only agent engaged in model averaging. Individual forecasters recursively update a single model. However, forecasters know their forecasts are being averaged by the policymaker, and unlike the policymaker, they are aware of the potential feedback from forecasts to prices. As a result, their model specifications

¹⁵See Sargent (1999, chpt. 6) for another example of this problem.

¹⁶Note, the unit root in the perceived model in eq. (6.33) implies that its spectral density is not well defined. (It is infinite at $\omega = 0$). In the numerical calculations, we approximate by setting $(1-L) = (1-\eta L)$, where $\eta = .995$. This means that our frequency domain objective is ill-equipped to find the degenerate fixed point where $\gamma = \Gamma = 0$. When this is the case, the true model exhibits i.i.d fluctuations around a mean of $\delta/(1-\alpha)$, while the agent’s perceived model exhibits i.i.d fluctuations around a mean of zero. The only difference between these two processes occurs at frequency zero, which is only being approximated here.

Figure 9: Best Response Mapping $\gamma = B(\Gamma)$

must include beliefs about how other forecasters are forecasting. These beliefs take the form of conjectures about other forecasters' beliefs about stationarity. \mathcal{M}_1 thinks \mathcal{M}_0 uses a constant parameter model, while \mathcal{M}_0 thinks \mathcal{M}_1 uses a random coefficients model with i.i.d. fluctuations around a constant mean. Both beliefs are misspecified, since they neglect the other agent's learning dynamics. However, as in traditional (representative agent) learning models, these misspecifications disappear asymptotically.

To be more specific, we can summarize the beliefs of \mathcal{M}_0 and \mathcal{M}_1 as a pair of perceived state space models. The perceived observation equation for \mathcal{M}_0 is

$$p_t = (1 - \pi_t)\beta_t(0)z_t + \pi_t\bar{\beta}_t(1)z_t + \sigma\epsilon_t \quad (6.34)$$

while the perceived observation equation for \mathcal{M}_1 is

$$p_t = (1 - \pi_t)\bar{\beta}_t(0)z_t + \pi_t\beta_t(1)z_t + \sigma\epsilon_t \quad (6.35)$$

Although forecasters are assumed to know the policymaker's current model weight, π_t , the coefficients on z_t are assumed to be hidden states. An overbar is used to represent an agent's belief about the other agent's model.

\mathcal{M}_0 has the following perceived state transition equation

$$\begin{bmatrix} \beta_t(0) \\ \bar{\beta}_t(1) \end{bmatrix} = \begin{bmatrix} \beta_{t-1}(0) \\ \bar{\beta}_{t-1}(1) \end{bmatrix} + \begin{bmatrix} 0 \\ \epsilon_{v,t} \end{bmatrix} \quad (6.36)$$

where $\epsilon_{v,t}$ is i.i.d Gaussian white noise. That is, \mathcal{M}_0 thinks the coefficient on z_t 'should be' constant, but at the same time recognizes the influence of a competing forecast based on alternative assumptions. Both equations are misspecified, since they both neglect the presence of learning dynamics. Misspecification about $\beta_t(0)$ will disappear at the usual t^{-1} rate. Misspecification about $\bar{\beta}_t(1)$ will persist, since in fact \mathcal{M}_1 's parameter estimate is governed by the Ornstein-Uhlenbeck process described in Proposition 2.1. However, viewed on the appropriate time-scale, this misspecification will also disappear. From the viewpoint of \mathcal{M}_0 , $\beta_t(0)$ appears to be drawn from the stationary distribution which is a solution of the Ornstein-Uhlenbeck process, centered at a constant.

\mathcal{M}_1 has the following perceived state transition equation

$$\begin{bmatrix} \bar{\beta}_t(0) \\ \beta_t(1) \end{bmatrix} = \begin{bmatrix} \bar{\beta}_{t-1}(0) \\ \beta_{t-1}(1) \end{bmatrix} + \begin{bmatrix} 0 \\ \sigma_v v_t \end{bmatrix}. \quad (6.37)$$

That is, \mathcal{M}_1 thinks the coefficient on z_t ‘should be’ drifting, but at the same time thinks rival forecasters are using a constant parameter model. Once again, these equations are misspecified during the transition, but will be confirmed in the limit.

Both forecasters estimate their perceived state space models using the Kalman filter. This generates a sequence of hidden state estimates which can then be substituted into the perceived observation equations to generate sequences of price forecasts.

$$\begin{aligned} \hat{p}_t(0) &= [(1 - \pi_t)\beta_{t+1|t}(0) + \pi_t\bar{\beta}_{t+1|t}(1)]\rho z_t \\ \hat{p}_t(1) &= [(1 - \pi_t)\bar{\beta}_{t+1|t}(0) + \pi_t\beta_{t+1|t}(1)]\rho z_t \end{aligned} \quad (6.38)$$

The policymaker takes these two forecasts and averages them using the known weight,

$$\hat{p}_t = (1 - \pi_t)\hat{p}_t(0) + \pi_t\hat{p}_t(1) \quad (6.39)$$

The actual time- t price is then determined by the (unknown) structural model in eq. (2.1),

$$p_t = \delta z_t + \alpha \hat{p}_t + \sigma \epsilon_t \quad (6.40)$$

Notice that if equations (6.38) are substituted into (6.39), which is then substituted into the actual price equation (6.40), we observe that the agents’ models in (6.34) and (6.35) suffer from an additional form of misspecification, since each forecaster fails to recognize that their own forecast embodies a form of model averaging in its response to the other forecasters’ forecasts. However, like the neglected learning dynamics, this misspecification will disappear in the limit.

We can follow virtually the same analysis as in Sections 5. We state the main result without proof.

Proposition 6.2. *In the decentralized equilibrium featuring three players, we again have $\lim_{t \rightarrow \infty} \beta_t(1) = \frac{\delta}{1 - \alpha\rho}$ weakly, while $\lim_{t \rightarrow \infty} \beta_t(0) = \frac{\delta}{1 - \alpha\rho}$ with probability 1. $\pi_t \rightarrow \{0, 1\}$ with probability 1, as $t \rightarrow \infty$. $\forall \varepsilon > 0$,*

$$\lim_{\sigma_v^2 \rightarrow 0} \lim_{T \rightarrow \infty} \mathbf{E} \frac{T_1}{T} = 1$$

where T_1 is the number of periods π_t is within $\varepsilon > 0$ neighborhood of 1.

It must be emphasized that the two forecasters have different views about the stationarity of the economy, even in the long run. Since \mathcal{M}_1 believes that $\beta_t(1)$ evolves according to a random walk, he expects to see the coefficient drift. If feedback is strong and fundamentals are persistent, this belief can persist for a *very* long time, as noted in the discussion following Proposition 2.1. Interestingly, given knowledge of the policymaker’s model averaging strategy, the persistence of his misspecification is independent of π_t . If π_t is close to 1, his belief will be confirmed for as long as it takes to reject a near random walk. On the other hand, if π_t is close to 0, the conditional variance of prices is close to σ^2 . But, this small variation of the price does not contradict his belief that $\beta_t(1)$ evolves according to a random walk, because \mathcal{M}_1 realizes that the actual price is partly determined by π_t , which is selected by the policymaker. Hence, he can rationalize the small gap between the actual variance of price and σ^2 coherently, as a consequence of π_t being close to 0, rather than as invalidating his belief that the coefficient would drift if his model were given more weight.

The fact that \mathcal{M}_1 maintains his belief even in the neighborhood of $\pi_t = 0$ is crucial for the escape dynamics of π_t from the neighborhood of 0. When the \mathcal{M}_0 forecast fares badly due to

a sequence of unusual events, the \mathcal{M}_1 forecast can quickly adjust to the event, which triggers a self-reinforcing escape of π_t away from 0 to 1.

At the same time, \mathcal{M}_0 maintains his belief that the economy is stationary. If π_t is close to 0, his belief is confirmed by the data, since the data is generated by what \mathcal{M}_0 forecasts. If π_t is close to 1, the data is mostly generated by what \mathcal{M}_1 forecasts. However, \mathcal{M}_0 infers that $\beta_t(1)$ is moving around $\delta/(1 - \alpha\rho)$, and can rationalize the gap between observed price variation and σ^2 in terms of the noise $\epsilon_{v,t}$ introduced by \mathcal{M}_1 forecaster. Hence, the data are also consistent with what \mathcal{M}_0 believes.

This feature of persistently misspecified beliefs is reminiscent of the Berk-Nash equilibrium of Esponda and Pouzo (2015), who distinguish between Nash and self-confirming equilibria. In a Nash equilibrium, each player's belief must be correctly specified, and their strategies must be econometrically identified. If we relax the requirement that equilibrium strategies be identified, we have a self-confirming equilibrium. If we also relax the requirement of correctly specified models, we then have the Berk-Nash equilibrium of Esponda and Pouzo (2015). Our model is not game theoretic, but nonetheless captures key features of a Berk-Nash equilibrium. We formulate the beliefs of each player in terms of a perceived law of motion. Since the perceived law of motion of \mathcal{M}_0 about $\beta_t(1)$ is not equal to the actual law of motion of $\beta_t(1)$, the belief of \mathcal{M}_0 is misspecified. In the long run, \mathcal{M}_0 cannot identify whether $\beta_t(1)$ is evolving according to a random walk, or whether it is drawn from a stationary distribution. As a result, \mathcal{M}_0 maintains a misspecified belief even in the long run, different from what \mathcal{M}_1 perceives about the economy.

7. CONCLUSION

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

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

APPENDIX A. PROOF OF PROPOSITION 4.2

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. (4.16) 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

$$\begin{aligned}\pi_{t+1} &\leq \pi_t + \pi_t(1 - \pi_t)M \\ &< \pi_t + \pi_t(1 - \pi_t) \\ &\leq 1\end{aligned}$$

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

APPENDIX B. PROOF OF LEMMA 4.3

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

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

since $\{\pi_t\} \in \Pi_0$. Depending upon the rate of convergence (or the time scale according to which π_t converges to π^*), we have to treat π_t has already converged to π^* .¹⁷

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

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

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

Under the assumption of Gaussian distributions,

$$\ln \frac{A_t(0)}{A_t(1)} = -\frac{(p_t - \rho\beta_t(0)z_t)^2}{2(\sigma^2 + \Sigma_t(0)z_t^2)} + \frac{(p_t - \rho\beta_t(1)z_t)^2}{2(\sigma^2 + \Sigma_t(1)z_t^2)} + \frac{1}{2} \ln \left[\frac{\sigma^2 + \Sigma_t(1)z_t^2}{\sigma^2 + \Sigma_t(0)z_t^2} \right]. \quad (\text{B.41})$$

Since the first two terms are normalized Gaussian variables,

$$\mathbb{E} \ln \frac{A_t(0)}{A_t(1)} = \mathbb{E} \frac{1}{2} \ln \left[\frac{\sigma^2 + \Sigma_t(1)z_t^2}{\sigma^2 + \Sigma_t(0)z_t^2} \right].$$

Recall (2.6), and note that $\Sigma_t(0) \rightarrow 0$. On the other hand, $\Sigma_t(1)$ is uniformly bounded away from 0, as $t \rightarrow \infty$, and the lower bound converges to 0, as $\sigma_v^2 \rightarrow 0$. Thus, $\beta_t(1)$ evolves on a faster time scale than $\beta_t(0)$. In calculating the limit value of (B.41), we first let $\beta_t(1)$ reach its own ‘‘limit’’, and then let $\beta_t(0)$ go to its own limit point.

Let $p_t^e(i)$ be the period- t price forecast by model i ,

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

Since

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

the forecast error of model 1 is

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

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

$$\lim_{t \rightarrow \infty} \mathbb{E} [\alpha\rho(1 - \pi_t)\beta_t(0) + (\alpha\rho\pi_t - 1)\beta_t(1) + \delta] = 0$$

in any limit point of the Bayesian learning dynamics.¹⁸ Since $\beta_t(1)$ evolves at a faster rate than $\beta_t(0)$, we can treat $\beta_t(0)$ as a constant. Since $\pi_t = \pi^*$, we treat π_t as constant also.¹⁹ Define

$$\bar{\beta}(1) = \lim_{t \rightarrow 0} \mathbb{E}\beta_t(1)$$

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

$$\lim_{\Sigma_t(1)} \lim_{t \rightarrow 0} [\alpha\rho(1 - \pi_t)\beta_t(0) + (\alpha\rho\pi_t - 1)\bar{\beta}(1) + \delta] + \mathbb{E}(\alpha\rho\pi_t - 1)(\beta_t(1) - \bar{\beta}(1)) = 0.$$

¹⁷If 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.

¹⁸Existence is implied by the tightness of the underlying space.

¹⁹If π_t evolves on a slower time scale than $\beta_t(1)$, we treat π_t as a constant, while investigating the asymptotic properties of $\beta_t(1)$.

Thus,

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

Define the deviation from the long-run mean as

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

Model 1's mean-squared forecast error is then

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

Note that

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

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

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

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

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

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

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

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

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

Thus, in the long run

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

in distribution, as $\Sigma_t(1) \rightarrow 0$ or equivalently, $\sigma_v^2 \rightarrow 0$. Note that

$$\lim_{t \rightarrow \infty} \frac{\mathbb{E}(p_t - p_t^e(0))^2}{\mathbb{E}(p_t - p_t^e(1))^2} > 1 \tag{B.42}$$

if and only if

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

Now, notice that

$$\frac{\alpha\rho\pi_t}{1 - \alpha\rho\pi_t} < 1$$

if and only if

$$\alpha\rho\pi_t < \frac{1}{2}.$$

Hence, if (B.42) holds for some $t \geq 1$, then it holds again for $t + 1$, and vice versa. Thus, π_t continues to increase or decrease, if the inequality holds in either direction. Recall that $\pi^* = \lim_{t \rightarrow \infty} \pi_t$. Convergence to π^* can occur only if (B.42) holds with equality for all $t \geq 1$, which is a zero probability event. We conclude that $\pi^* \in (0, 1)$ occurs with probability 0. \square

APPENDIX C. PROOF OF LEMMA 4.4

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 $t \rightarrow \infty$, we know $\phi_t \rightarrow \phi^*$ with probability 1. Hence, we have

$$\begin{aligned} \lim_{t \rightarrow \infty} t^\alpha(\pi_t - \pi_{t-1}) &= \lim_{t \rightarrow \infty} \frac{t^\alpha(e^{-\phi^*} - 1)e^{t\phi^*}}{(1 + e^{t\phi^*})(1 + e^{(t-1)\phi^*})} \\ &= (e^{-\phi^*} - 1) \lim_{t \rightarrow \infty} \frac{t^\alpha}{(1 + e^{-t\phi^*})(1 + e^{t\phi^*}e^{-\phi^*})} \end{aligned}$$

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

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

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

APPENDIX D. PROOF OF PROPOSITION 5.3

D.1. Preliminaries. Lemma 4.3 and Lemma 4.4 shows that $\beta_t(1)$ moves at the fastest time scale, followed by $\beta_t(0)$ and then π_t . The same reasoning also shows the domain of attraction for $\pi = 0$ is

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

and the domain of attraction for $\pi = 1$ is

$$\mathcal{D}_1 = \left\{ (\pi, \beta(0), \beta(1)) \mid \mathbb{E} \log \frac{A_t(0)}{A_t(1)} < 0 \right\}.$$

Since $\beta_t(1)$ does not trigger the escape from one domain of attraction to another, let us focus on $(\pi, \beta(0))$, assuming that we are moving according to the time scale of $\beta_t(0)$. A simple calculation shows that \mathcal{D}_0 has a narrow symmetric shape of $(\pi, \beta(0))$, centered around

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

with the base

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

along the line $\pi = 0$ where

$$d = \frac{\sqrt{\bar{\Sigma}}}{1 - \alpha\rho}. \tag{D.43}$$

Note that since $\bar{\Sigma} \rightarrow 0$ as $\sigma_v \rightarrow 0$,

$$\lim_{\sigma_v \rightarrow 0} d = 0.$$

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 $\beta_t(0), \beta_t(1) \rightarrow \frac{\delta}{1 - \alpha\rho}$,

$$\phi^* = \mathbb{E} \log \frac{A_t(0)}{A_t(1)} = \mathbb{E} \frac{1}{2} \log \frac{\text{MSE}(1)}{\text{MSE}(0)}$$

is defined for $\beta_t(0) = \beta_t(1) = \frac{\delta}{1 - \alpha\rho}$, and $\pi = 1$ or 0 .

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

$$\phi_-^* < 0 < \phi_+^*$$

and

$$\phi_-^* + \phi_+^* > 0.$$

Recall r_0 and r_1 are the rate functions of \mathcal{D}_0 and \mathcal{D}_1 . For fixed $\sigma_v > 0$, define

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

and

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

Then,

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

Our goal is to show that $\exists \bar{\sigma}_v > 0$ such that

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

D.2. Escape probability from \mathcal{D}_1 . Consider a subset of \mathcal{D}_1

$$\mathcal{D}'_1 = \{(\beta(1), \beta(0), \pi) \mid \pi > \bar{\pi}\}.$$

For fixed $\sigma_v > 0$, define

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

and

$$r_1^* = \liminf_{\sigma_v \rightarrow 0} r_1^*(\sigma_v).$$

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

$$\phi_t > 0.$$

We know

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log \mathbf{P}(\exists t, \phi_t > 0 \mid \phi_1 = \phi_-^*) = r_1^*(\sigma_v).$$

We claim that

$$r_1^* > 0. \tag{D.44}$$

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

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

which implies $\bar{\Sigma} \rightarrow 0$. Note that

$$\phi_t > 0$$

if and only if

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

if and only if

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

if and only if

$$\frac{1}{t} \sum_{k=1}^t \left[\frac{\log \frac{A_t(0)}{A_t(1)} - \mathbf{E} \log \frac{A_t(0)}{A_t(1)}}{\bar{\Sigma}} \right] > -\frac{\phi_-^*}{\bar{\Sigma}}. \tag{D.45}$$

A straightforward calculation shows

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

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

$$\frac{\log \frac{A_t(0)}{A_t(1)} - \mathbf{E} \log \frac{A_t(0)}{A_t(1)}}{\bar{\Sigma}}$$

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

We can write

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

where

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

and

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

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

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

A simple calculation shows

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

Since $\Sigma_t(1) \rightarrow \bar{\Sigma} > 0$, 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 $\forall t \geq T(\epsilon)$,

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

Thus, as $t \rightarrow \infty$,

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

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

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

is precisely

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

That means, the large deviation properties of the left hand side of (D.45) is the same as the large deviation properties of Z_t . Since Z_t is the sample average of a martingale difference, a standard argument from large deviation theory implies that its rate function is strictly positive for given $\sigma_v > 0$. We normalized the martingale difference by dividing each term by $\bar{\Sigma}$ so that the second moment of

$$\frac{\log \frac{A_t(0)}{A_t(1)} - \mathbf{E}_t \log \frac{A_t(0)}{A_t(1)}}{\bar{\Sigma}}$$

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

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

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

$$\lim_{t \rightarrow \infty} \log \mathbf{P} \left(\frac{1}{t} \sum_{k=1}^t \left[\frac{\log \frac{A_t(0)}{A_t(1)} - \mathbf{E} \log \frac{A_t(0)}{A_t(1)}}{\bar{\Sigma}} \right] \geq -\frac{\phi_-^*}{\bar{\Sigma}} \right) = \lim_{t \rightarrow \infty} \log \mathbf{P} \left(Z_t \geq -\frac{\phi_-^*}{\bar{\Sigma}} \right) = r_1^*(\sigma_v) \quad (\text{D.46})$$

and

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

as desired.

D.3. Escape probability from \mathcal{D}_0 . Recall that $\beta_t(0)$ evolves according to

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

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

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

and

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

D.4. Conclusion. Recall (D.43) to notice that

$$\lim_{\sigma_v \rightarrow 0} d = 0.$$

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

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

Observe

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

since the exit occurs at the most likely exit point at which r_0 is determined, while $r_0^*(d)$ is determined at a particular exit point. Since $\mathcal{D}'_1 \subset \mathcal{D}_1$,

$$r_1^*(\sigma_v) \leq r_1(\sigma_v).$$

Thus, for any $\sigma_v > 0$ sufficiently close to 0,

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

from which

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

follows.

REFERENCES

- ACEMOGLU, D., V. CHERNOZHUKOV, AND M. YILDIZ (2016): “Fragility of Asymptotic Agreement Under Bayesian Learning,” *Theoretical Economics*, 11, 187–225.
- 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.

- BORKAR, V. S. (2008): *Stochastic Approximation: A Dynamical Systems Viewpoint*. Cambridge University Press.
- BRAY, M., AND N. SAVIN (1986): "Rational Expectations Equilibria, Learning, and Model Specification," *Econometrica*, 54, 1129–60.
- BULLARD, J. (1992): "Time-Varying Parameters and Nonconvergence to Rational Expectations under Least-Squares Learning," *Economics Letters*, 40, 159–66.
- 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.
- CHO, I.-K., N. WILLIAMS, AND T. J. SARGENT (2002): "Escaping Nash Inflation," *Review of Economic Studies*, 69, 1–40.
- DEMO, 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 (2015): "Berk-Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models," Washington University and UC Berkeley.
- EVANS, G. W., AND S. HONKAPOHJA (1993): "Adaptive Forecasts, Hysteresis and Endogenous Fluctuations," *Economic Review of Federal Reserve Bank of San Francisco*, 1, 3–13.
- (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press.
- EVANS, G. W., S. HONKAPOHJA, T. J. SARGENT, AND N. WILLIAMS (2013): "Bayesian Model Averaging, Learning, and Model Selection," in *Macroeconomics at the Service of Public Policy*, ed. by T. J. Sargent, and J. Vilmunen. Oxford University Press.
- HANSEN, L. P. (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.
- 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.
- 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.
- MCGOUGH, B. (2003): "Statistical Learning with Time-Varying Parameters," *Macroeconomic Dynamics*, 7, 119–39.
- 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.
- 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.
- SARGENT, T. J., AND N. WILLIAMS (2005): "Impacts of Priors on Convergence and Escapes from Nash Inflation," *Review of Economic Dynamics*, 8, 360–91.
- 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.

DEPARTMENT OF ECONOMICS, UNIVERSITY OF ILLINOIS, 1407 WEST GREGORY, URBANA IL 61801 USA

E-mail address: inkoocho@illinois.edu

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

DEPARTMENT OF ECONOMICS, SIMON FRASER UNIVERSITY, 8888 UNIVERSITY DRIVE, BURNABY, B.C. CANADA. V5A 1S6

E-mail address: kkasa@sfu.ca

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