COMPETING MODELS WITH FEEDBACK

IN-KOO CHO AND KENNETH KASA

Abstract. This paper studies misspecified Bayesian learning with endogenous data when models compete via model averaging. A decisionmaker forecasts an endogenous price sequence by averaging the forecasts from two models, one with constant parameters and one with time-varying parameters. The time-varying parameters model is misspecified. It excludes a relevant explanatory variable and fails to recognize the presence of model averaging. In contrast, the constant parameters model includes all relevant explanatory variables and recognizes that prices are generated by averaging the two forecasts. If expectational feedback is weak, the correctly specified constant parameters model prevails. However, if feedback is strong and the excluded fundamentals are not too important, the under-parameterized time-varying parameters model survives the competition, and prices become endogenously nonstationary. Simple time-varying parameter models do well because they better respond to time variation in the data that their own use generates.

JEL Classification Numbers: C52, D80

What the government observes depends partly on what it believes. – Sargent (1999, p. 75)

1. Introduction

Model use and construction in economics confront four key challenges: (1) Uncertainty, (2) Misspecification, (3) Feedback, and (4) Competition. In principle, dealing with model uncertainty is straightforward: specify a prior and apply Bayes Rule. In practice, it is not so easy. Relevant parameter spaces are infinite dimensional, so priors are almost surely misspecified. Fortunately, the effects of prior misspecification are well understood. Berk (1966) shows that misspecified Bayesian learning converges to a model within the prior that minimizes the Kullback-Leibler distance to the unknown true model.

The most serious challenge is *feedback*. The data are endogenous and react to the forecaster's own beliefs. By itself, feedback is not a problem. Dealing with feedback was what the Rational Expectations revolution was all about. However, the Rational Expectations hypothesis abstracts from model uncertainty and misspecification. In response, researchers have begun studying misspecified learning with feedback in the past decade. Esponda and Pouzo (2016) propose the concept of a Berk-Nash equilibrium. In a Berk-Nash equilibrium, players use optimal strategies given their own beliefs. At the same time, each player's beliefs minimize an *endogenous* Kullback-Leibler distance to the true model. Esponda and Pouzo (2021) and Fudenberg, Lanzani, and Strack (2021) extend and refine the Berk-Nash equilibrium concept.

Our paper studies misspecified Bayesian learning with feedback but with an added twist. Here, we assume models must *compete* via model averaging. This setting is particularly

Date: March 16, 2025.

relevant for macroeconomic policy. Blinder (1998, p. 12) notes that model averaging is used at the Fed, albeit informally:

> *Parameter uncertainty, while difficult, is at least a relatively well defined problem. Selecting the right model from among a variety of non-nested alternatives is another matter entirely... My approach to this problem while on the Federal Reserve Board was relatively simple: Use a wide variety of models and don't ever trust any one of them too much....My usual procedure was to simulate a policy on as many of these models as possible, throw out the outliers, and average the rest to get a point estimate of a dynamic multiplier path.*

Despite its practical importance, surprisingly little work has been done on how model competition interacts with misspecification and feedback.

Olea, Ortoleva, Pai, and Prat (2022) study model competition without feedback. They consider an environment where a collection of (linear) models compete against each other based on subjective mean-squared prediction error. Using a standard bias/variance decomposition of mean-squared error, they show that when data are scarce, small parsimonious models prevail. However, as data become more plentiful, the prevailing model becomes larger. Interestingly, they show that an overparameterized model has a positive probability of prevailing, even in the limit.

Evans, Honkapohja, Sargent, and Williams (2013) and Cho and Kasa (2017) study model competition with feedback, but in a setting where models only differ in their assumptions about parameter stability. They consider a policymaker who must forecast a future endogenous price. He hires two competing forecasters to construct models and report forecasts. The policymaker forms his forecast as a recursively updated weighted average of the two forecasts. The true data-generating process is unknown and relates the current price to the policymaker's forecast of the next period's price and an exogenous fundamental. One forecaster thinks the environment is stationary and estimates (recursively) a constant parameters (CP) model. The other forecaster thinks the environment is nonstationary and estimates a time-varying parameters (TVP) model. Importantly, both forecasters have misspecified models of endogeneity. Since the coefficients are, in fact, constant, there is a sense in which the CP forecaster has the correct model. In a Rational Expectations equilibrium, the parameters *are* constant. His only misspecification is that he neglects the transitional learning dynamics. If he operated in isolation, these transitional dynamics would dissipate over time as long as standard E-stability conditions are satisfied (Evans and Honkapohja (2001)). However, Cho and Kasa (2017) show that if the CP forecaster must compete in real-time against the TVP forecaster via model averaging, the TVP forecaster will ultimately prevail. As a result, the economy is more unstable than it otherwise would be. Hence, in this setting, competition is *undesirable*. Competition is great when the rules are fixed and fair. However, with feedback, the TVP model can effectively alter the rules of the game in its favor. This is because the TVP model can better adapt to feedback-induced non-stationarity in the data.

This paper revisits the analysis in Cho and Kasa (2017) by giving the CP model two advantages. First, the TVP model is under-parameterized. It excludes the exogenous fundamentals. We show that as long as the variance of the omitted fundamental is not too

large, the TVP model still dominates. Its ability to react to feedback more than offsets the fact that it is under-parameterized. This strengthens the results in Cho and Kasa (2017) and potentially explains why, in practice, TVP models are often relatively parsimonious.[1]

Of course, despite including all relevant explanatory variables, the CP model is at a serious disadvantage since it fails to recognize that its forecast must compete with the TVP forecast. Even if the CP forecaster is convinced the environment is stationary, he should realize that the policymaker's averaging strategy will indirectly induce non-stationarity. Thus, our second advantage to the CP model is that it allows the CP forecaster to be aware of the competition. He effectively imports some non-stationarity into his forecast by averaging his forecast with that from the TVP model. The forecasting rule of this "augmented" CP model is correctly specified according to Esponda and Pouzo (2016). Since the CP forecaster cannot observe the probability weight assigned to each model by the decision maker, the CP forecaster estimates the weight using recursive least squares.

Surprisingly, we find that even now, the TVP can compete successfully. However, since the CP forecaster is increasingly copying the TVP model as the weight on the TVP model increases, it is no longer the case that the TVP model dominates in the sense that its weight converges to 1. Instead, we obtain a strictly interior long-run equilibrium, where the superior adaptability of the TVP model balances its underparameterization. As the strength of feedback increases, the long-run equilibrium weight on the TVP model converges to 1.

These results are reminiscent of examples presented by Blume and Easley (1982) and Bray and Kreps (1987). They show that even if an agent's prior contains the Rational Expectations equilibrium model, Bayesian updating may not converge to it if it does not adequately describe out-of-equilibrium transition dynamics. Doesn't this just mean we need to extend the prior to include these transition dynamics? In principle, yes, but as discussed in Bray and Kreps (1987) and Nachbar (1997), this is a methodological blind alley that we prefer to avoid. This is particularly true in the multiple-agents/macroeconomic policy settings. Our goal is more descriptive and more practical. We want to provide a plausible explanation for parameter instability in applied time-series econometrics. Specifying complex priors that fully capture both feedback and averaging among multiple models and agents would counter the spirit of such an analysis. Instead, the simple augmented CP model provides a useful laboratory for studying how feedback and model competition interact to determine which models are used in practice.

The joint dynamics of the data, model coefficients, and model weights is a high-dimensional, nonlinear, stochastic process. The key to making the analysis tractable is to exploit the fact that different processes evolve on different time scales. This effectively reduces the dimensionality of the problem by allowing us to study the interactions among smaller dimensional subsystems. Marcet and Sargent (1989) were the first to import time-scale separation methods into the economics literature. Our problem is more complex than the one studied by Marcet and Sargent (1989), which just exploited a time-scale

[1]Our result is related to a classic paper by Nelson (1972), who showed that simple univariate ARIMA models produced superior forecasts to the large FRB-MIT-Penn model. It is also related to the literature on Bayesian VARs, where it was discovered that the Minnesota prior is an effective way of containing estimation variance. (The Minnesota prior postulates a nondogmatic diagonal VAR coefficient matrix with random walk parameter drift. See, e.g., Doan, Litterman, and Sims (1984)).

separation between the data and the coefficients of a single model. Here, multiple models evolve on different time scales and model weights that evolve on a different time scale. Nonetheless, by taking limits in the appropriate order, we can study the dynamics of a *sequence* of Ordinary Differential Equations (Borkar (2008)).

The remainder of the paper is organized as follows. Section 2 presents the data-generating process and the two competing models and discusses how each model is updated using the Kalman filter. Section 3 discusses model averaging. We consider two cases. We start with a simple case where the CP forecaster is unaware of the policymaker's averaging process. We then augment the CP forecasting rule by incorporating the policymaker's averaging process into the forecasting rule. The CP forecaster estimates the probability weight assigned to the TVP model using recursive least squares. Section 4 illustrates the main results. We show that if feedback is strong enough and the excluded fundamental is not too important, then convergence to a unique, stable, steady state of the probability weight occurs, converging to one as feedback strengthens. Even though the augmented CP model is correctly specified in the sense of Esponda and Pouzo (2016), the under-parameterized TVP model generates the data with a probability close to 1. Section 5 presents results from the numerical exercises that illustrate our theoretical results. This is useful since our results are asymptotic and predict what happens as shock variances tend to zero in the appropriate order. The simulations show that convergence occurs continuously with respect to the parameter values. Finally, Section 6 offers concluding remarks, and a technical appendix contains the proofs of our key convergence results.

2. Description

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

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

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

$$f_t = \rho f_{t-1} + \sigma_f \epsilon_{f,t} \tag{2.2}$$

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

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

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

$$p_t(0) = \beta_t(0) + \frac{\delta}{1-\alpha\rho} f_t + \sigma \epsilon_t \tag{2.4}$$

$$\beta_t = \beta \tag{2.5}$$

for some β. To simplify the ensuing analysis, we assume the agent knows the coefficient on f_t, and must only estimate the constant term β. Requiring the agent to estimate the coefficient on f_t as well would strengthen our result since we shall see that, ultimately, this model is dominated by an alternative model, which is simpler but under-parameterized. Note that without model competition, the Rational Expectations equilibrium is a special case of this agent's perceived model, with

$$\beta = \frac{\gamma}{1-\alpha}.$$

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

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

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

$$\beta_{t+1}(0) = \beta_t(0) + \left(\frac{\Sigma_t(0)}{\sigma^2 + \Sigma_t(0)}\right)\left(p_t - \beta_t(0) - \frac{\delta}{1-\alpha\rho} f_t\right) \tag{2.6}$$

$$\Sigma_{t+1}(0) = \Sigma_t(0) - \frac{(\Sigma_t(0))^2}{\sigma^2 + \Sigma_t(0)}. \tag{2.7}$$

Following standard analysis (Evans and Honkapohja (2001)), we can prove the convergence and the stability properties of the learning dynamics under specification $\mathcal{M}_0$.

Proposition 2.1. *Given $\alpha \in (0,1)$ and the belief that parameters are constant, $\Sigma_t(0)$ converges to zero at rate t^{-1}, and*

$$\beta_t(0) \to \frac{\gamma}{1-\alpha}$$

with probability 1.

2.3. **Learning with a simpler but misspecified model.** Now, suppose there is a second agent with different prior beliefs, described by the state space model:

$$p_t(1) = \beta_t(1) + \sigma\epsilon_t \tag{2.8}$$

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

where v_t is standard Gaussian white noise, which is orthogonal to all other variables. The innovation variance σ_v^2 reflects the agent's priors about parameter drift. The larger it is, the more aggressively the agent will revise his beliefs in response to forecast errors. We treat it as a fixed parameter.[2]

A potentially serious specification error here is that the agent's model excludes the fundamental f_t. The agent entertains a time-varying parameters (TVP) model to capture variation from the missing variable.

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

The missing variable f_t produces a gap between the perceived law of motion and the actual law of motion of p_t. Let $\beta_t(1)$ and $\Sigma_t(1)$ be the mean and the variance of the posterior

[2]Section 5 of Cho and Kasa (2017) shows that our results extend to the case where σ_v^2 is estimated.

distribution about $\beta_t(1)$ conditioned on information at $t-1$, which is computed from a Gaussian prior. The evolution of $\beta_t(1)$ and $\Sigma_t(1)$ is dictated by the new Kalman filter:

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

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

while the actual price p_t is driven by (2.1)

$$p_t = \gamma + \alpha\beta_t(1) + \delta f_t + \sigma\epsilon_t.$$

Since the Kalman gain no longer vanishes, we cannot obtain convergence with probability 1. Instead, we obtain convergence in distribution.

Proposition 2.2. *As $\sigma_v^2 \to 0$, $\beta_t(1)$ converges weakly to the solution of the following diffusion process*

$$d\beta = (1-\alpha)\left[\frac{\gamma}{1-\alpha} - \beta\right]d\tau + \sigma_v\sqrt{1 + \frac{\sigma_f^2}{\sigma^2(1-\rho^2)}}dW_\tau \tag{2.12}$$

where dW_τ is a standard Wiener process.

Proof. See Kushner and Yin (1997). □

Notice that (2.12) differs from the perceived law in (2.9) in two ways. First, it is not a random walk but exhibits mean reversion. This inconsistency reflects the agent's neglect of the effects induced by his own learning. However, it will be quite small and difficult to detect when α is close to 1, and the innovation variance is small. Second, the prior volatility of $\beta_t(1)$ in (2.9) will be less than the steady state posterior volatility in (2.12) by a factor of $1 + \sigma_f^2/[\sigma^2(1-\rho^2)]$. This reflects the fact that the agent uses parameter drift to capture the dynamics of the omitted variable f_t. This inconsistency can be eliminated by allowing the agent to estimate σ_v.

3. Model averaging

Model competition can take many forms. Olea, Ortoleva, Pai, and Prat (2022) motivate their analysis in an auction setting. However, when studying competition with feedback, the most natural setting is model averaging. A decisionmaker mixes the predictions from two models to hedge his bets. Bayes rule provides an algorithm for updating the weights in response to new data.

Following Evans, Honkapohja, Sargent, and Williams (2013) and Cho and Kasa (2017), we consider a decentralized environment in which there is a separation between model construction and model use. This is particularly descriptive of macroeconomic policy. Letting π_t denote the current probability assigned by the policymaker to $\mathcal{M}_1$ (the TVP model), the policymaker's time-t forecast is becomes

$$\mathsf{E}_t p_{t+1} = \hat{p}_t = \pi_t p_t(1) + (1-\pi_t)p_t(0)$$

where $p_t(i)$ is the time-t forecast reported from $\mathcal{M}_i$. Since the TVP forecaster dogmatically omits the fundamental from his model, the forecast from $\mathcal{M}_1$ is just

$$p_t(1) = \beta_t(1)$$

where $\beta_t(1)$ is given by (2.8) and (2.9). We consider two cases for $p_t(0)$ depending on whether the CP forecaster is aware of the policymaker's averaging. We can nest these two cases as follows

$$p_t(0) = (1 - \hat{\pi}_t)\left[\beta_t(0) + \frac{\delta}{1 - \alpha\rho} f_t\right] + \hat{\pi}_t \beta_t(1) \tag{3.13}$$

where $\hat{\pi}_t$ is the CP forecaster's current estimate of the policymaker's weight. It is given by the recursive least squares estimate

$$\hat{\pi}_t = \hat{\pi}_{t-1} + \frac{1}{t}\left(\pi_t - \hat{\pi}_{t-1}\right). \tag{3.14}$$

In Evans, Honkapohja, Sargent, and Williams (2013) and Cho and Kasa (2017), the CP forecaster dogmatically believes that $\pi_t = 0$, ignoring the fact that the policymaker averages the two forecasts. The CP model is correctly specified only if the CP model is used in isolation. If the CP model is competing against the TVP model, the CP model can be improved by incorporating the policymaker's model averaging process in (3.13). The case where the CP forecaster neglects averaging can be studied by simply setting $\hat{\pi}_t = 0$ $\forall t$. We call the forecasting algorithm with (3.14) the *augmented CP model*. Since the augmented CP model replicates the functional form of the data-generating process, it is correctly specified in the sense of Esponda and Pouzo (2016).

We assume the CP forecaster does not know the policymaker's current weight. Instead, he thinks it is an unknown constant, which he must estimate himself. This is in keeping with our assumption that model construction and use are decentralized and occur among agents. His assumption that the weight is constant will be correct in the limit but neglects transition dynamics. In particular, the CP forecaster overlooks feedback between his forecasting rule and the data-generating process, which makes π_t an endogenous variable instead of an unobserved constant.

The asymptotic properties of π_t are the focus of our analysis. If the augmented CP model generates the data, π_t must be close to 0. Similarly, if π_t is close to 1, the (under-parameterized) TVP model generates the data.

The strength of feedback is quantified by $\alpha\rho$, where α represents expectational feedback and ρ captures how persistent the fundamental is. We show that if feedback is weak, $\pi_t \to 0$. This finding is consistent with the fact that the underlying rational expectations equilibrium is stationary. The augmented CP model correctly assumes that the stochastic process is stationary and is, therefore, a correctly specified model.

On the other hand, if feedback is sufficiently strong, $\pi_t \to \overline{\pi} < 1$, where $\overline{\pi}$ converges to 1 as feedback becomes stronger. Even if the augmented CP model is correctly specified, the under-parameterized TVP model eventually generates the data with a probability close to 1. This is because the TVP model is better able to respond to the feedback-induced non-stationarity of the data-generating process.[3]

[3]Cho and Kasa (2015) investigate a recursive model validation process, and show that even if a correctly specified model is in the decisionmaker's prior, the decision maker may end up using a misspecified model most of the time in the long run. The large deviation rate function around a self-confirming equilibrium

At this point, it is useful to collect together the formulas that govern the evolution of the six endogenous stochastic processes: $(\pi_t, \hat{\pi}_t, \beta_t(0), \Sigma_t(0), \beta_t(1), \Sigma_t(1))$.

$$\beta_{t+1}(0) = \beta_t(0) + \left(\frac{\Sigma_t(0)}{\sigma^2 + \Sigma_t(0)}\right)(p_t - p_t(0)) \tag{3.15}$$

$$p_t(0) = (1-\hat{\pi}_t)\left[\beta_t(0) + \frac{\delta}{1-\alpha\rho} f_t\right] + \hat{\pi}_t \beta_t(1) \tag{3.16}$$

$$\hat{\pi}_t = \hat{\pi}_{t-1} + \frac{1}{t}(\pi_t - \hat{\pi}_{t-1}) \tag{3.17}$$

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

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

$$p_t(1) = \beta_t(1) \tag{3.20}$$

$$\Sigma_{t+1}(1) = \Sigma_t(1) - \frac{(\Sigma_t(1))^2}{\sigma^2 + \Sigma_t(1)} + \sigma_v^2 \tag{3.21}$$

$$\hat{p}_t = (1-\pi_t)p_t(0) + \pi_t p_t(1) \tag{3.22}$$

$$\begin{aligned} p_t &= \gamma + \alpha \mathsf{E}_t \hat{p}_{t+1} + \delta f_t + \sigma \epsilon_t \\ &= \gamma + \alpha(\pi_t + (1-\pi_t)\hat{\pi}_t)\beta_t(1) + \alpha(1-\pi_t)(1-\hat{\pi}_t)\beta_t(0) \\ &\quad + \left(\alpha(1-\pi_t)(1-\hat{\pi}_t)\frac{\delta\rho}{1-\alpha\rho} + \delta\right) f_t + \epsilon_t. \end{aligned} \tag{3.23}$$

$$\frac{1}{\pi_{t+1}} - 1 = \frac{A_{t+1}(0)}{A_{t+1}(1)}\left(\frac{1}{\pi_t} - 1\right) \tag{3.24}$$

where $A_t(i)$ is the policymaker's perceived likelihood function of $\mathcal{M}_i$

$$A_t(i) = \frac{1}{\sqrt{2\pi(\sigma^2 + \Sigma_t(i))}} \exp\left[-\frac{(p_t - p_t(i))^2}{2\pi(\sigma^2 + \Sigma_t(i))}\right] \tag{3.25}$$

Note that in applying Bayes rule (B.43), the policymaker is unaware that the CP modeler is responding to his own averaging efforts.

Equations (B.34)-(B.44) represent a high dimensional, nonlinear, stochastic system. We exploit a couple of key features to simplify the analysis dramatically. First, since the models themselves are linear and Gaussian, the forecaster's confidence, as given by the conditional variances, $\Sigma_t(i)$, evolve *deterministically* and *exogenously*. This would not be the case in general. Second, since $\Sigma_t(0)$ converges to 0 at rate t^{-1} whereas $\Sigma_t(1)$ converges to a strictly positive limit, the two model estimates, $(\beta_t(0), \beta_t(1))$, evolve on different time-scales. $\beta_t(1)$ evolves much faster than $\beta_t(0)$. This opens the door to a classical time-scale separation strategy, in which we can fix $\beta_t(0)$ at a constant when studying the dynamics of $\beta_t(1)$. This simplifies the analysis to study the interactions among smaller dimensional subsystems.

---

determines the duration time. If a correctly specified model does not have the largest large deviation rate function, then a misspecified model is used almost always in the long run.

4. Results

The mathematical details are contained in the Appendix. Here, we merely summarize the results. First, let us consider the case where the CP forecaster ignores the non-stationarity generated by the TVP model: $\hat{\pi}_t \equiv 0$ as in Evans, Honkapohja, Sargent, and Williams (2013) and Cho and Kasa (2017). We analyze the dynamic system of (3.15)-(3.25) while replacing (3.17) by $\hat{\pi}_1 = 0$ and $\hat{\pi}_t = \hat{\pi}_{t-1}$ so that the CP forecaster ignores the fact that the policy maker averages the two predictions.

One might think that since his model includes f_t while the TVP model mistakenly excludes it, the CP model would outperform the TVP model, and $\pi_t \to 0$. Although this would certainly be the case if the variance of f_t is large, this is not particularly interesting since such obvious misspecification would likely be easily detected. We are interested in smaller, more subtle misspecifications. Interestingly, our first result shows that when f_t is not too important, the TVP model will dominate despite its under-parameterization.

To formalize the notion of domination, for small $\epsilon > 0$, define the duration time

$$T^{\epsilon}(T, \beta, \pi) = \#\left\{t \leq T \mid (\beta_t, \pi) \in \mathcal{N}_{\epsilon}(\beta, \pi)\right\}$$

as the number of periods when β_t and the sample average of π_t stay in an ϵ-neighborhood of (β, π). Following the same logic as in Cho and Kasa (2017), we can show that π_t converges to a bi-modal distribution concentrated at $\{0, 1\}$. Moreover, the TVP model dominates the CP model even though the TVP model excludes the fundamental f_t, as long as its variance is small and feedback is sufficiently strong.

Proposition 4.1. *Consider the dynamic system (3.15)-(3.24) while replacing* (3.17) *by* $\hat{\pi}_1 = 0$ *and* $\hat{\pi}_t = \hat{\pi}_{t-1}$. *If* $\alpha\rho > \frac{1}{2}$, *then*

$$\lim_{\sigma_f \to 0} \lim_{\sigma_v \to 0} \lim_{T \to 0} \mathsf{E} \frac{T^{\epsilon}(T, \frac{\gamma}{1-\alpha}, 1)}{T} = 1.$$

Proof. See Appendix A. □

Proposition 4.1 is striking in that the decisionmaker believes in the misspecified TVP model whenever expectational feedback is sufficiently strong. However, one could argue that the CP model is even more misspecified because it ignores that the decisionmaker averages the two forecasts, thus ignoring the induced non-stationarity generated by the TVP forecast.

A natural response is to augment the CP forecast by assuming that $\hat{\pi}_t$ follows (3.17). That is, the CP forecaster is aware that the decisionmaker averages the two forecasts. Since the CP forecaster does not observe π_t, he estimates it using (3.17), reflecting the assumption that π_t is an unknown constant. The resulting forecasting rule is the augmented constant parameter (CP) model. It is important that the CP forecaster maintains the hypothesis that the underlying data-generating process is stationary. Non-stationarity is assumed to only arise indirectly via the TVP model.

The augmented CP forecasting rule is correctly specified according to Esponda and Pouzo (2016), as (3.17) accurately captures the functional form of the actual law of motion in (3.23). Interestingly, we now find that π_t no longer converges to 1 as in Proposition 4.1. As $\pi_t \to 1$, the augmented CP model increasingly mimics the forecast from the TVP model. Moreover, the augmented CP model includes the fundamental variable f_t,

which the TVP model omits. Instead, there exists a $\overline{\pi} \in (0,1)$ such that $\pi_t \to \overline{\pi}$, and $\lim_{\alpha\rho \to 1} \overline{\pi} = 1$, implying that as the feedback becomes stronger, the TVP model still dominates the augmented CP model.

Theorem 4.2. *Case 1: If $\alpha\rho < \frac{2+\sqrt{3}}{4}$, then $(\beta, \hat{\pi}) = \left(\frac{\gamma}{1-\alpha}, 0\right)$ is the unique stable stationary point of $(\beta_t(0), \hat{\pi}_t)$ in (3.15)-(3.24), and $\pi_t \to 0$ in probability.*
Case 2: If $\alpha\rho > \frac{2+\sqrt{3}}{4}$, then $(\beta_t(0), \hat{\pi}_t)$ generated by (3.15)-(3.24) has two locally stable stationary points: $(\beta, \hat{\pi}) = \left(\frac{\gamma}{1-\alpha}, 0\right)$ and $(\beta, \hat{\pi}) = \left(\frac{\gamma}{1-\alpha}, \bar{\pi}\right)$ where

$$\bar{\pi} = \frac{4\alpha\rho - 1 + \sqrt{1 - 16\alpha\rho(1-\alpha\rho)}}{4\alpha\rho} < 1. \tag{4.26}$$

$\forall \epsilon > 0$, define

$$T^{\epsilon}(T : \beta, \hat{\pi}) = \#\left\{t \leq T \mid (\beta_t, \hat{\pi}_t) \in \mathcal{N}_{\epsilon}(\beta, \hat{\pi})\right\}$$

as the number of periods $(\beta_t, \hat{\pi}_t)$ stays in an ϵ-neighborhood of $(\beta, \hat{\pi})$. The proportion of time that $(\beta_t, \hat{\pi}_t)$ stays in a small neighborhood of the second stable stationary point converges to 1 in the long run as $\sigma_f, \sigma_v \to 0$.

$$\lim_{\sigma_f \to 0} \lim_{\sigma_v \to 0} \lim_{T \to \infty} \mathsf{E} \frac{T^{\epsilon}\left(T : \frac{\gamma}{1-\alpha}, \bar{\pi}\right)}{T} = 1. \tag{4.27}$$

Proof. See Appendix B. □

In the first part of Theorem 4.2, $\pi_t \to 0$ in probability, as $\hat{\pi}_t \to 0$. However, in the second part, where $\alpha\rho$ is close to 1, $\hat{\pi}_t \to \bar{\pi} < 1$. Note that the threshold determining whether feedback is 'strong' is higher here than in Theorem 4.1. This is because the CP model can now indirectly track time variation in the data. Theorem 4.2 implies that π_t does not converge to a point in the $[0,1]$ interval for a fixed $\sigma_f, \sigma_b > 0$, even though the expected value of π_t converges to $\bar{\pi}$. The combination of the non-stationarity of $\beta_t(1)$ and π_t generates considerable price volatility, as the numerical simulations indicate.

Note that $\overline{\pi} \to 1$ as $\alpha\rho \to 1$. If feedback is extremely strong, the augmented CP forecaster accurately forecasts the actual price p_t by estimating $\pi_t = 0$ accurately through $\hat{\pi}_t = 0$. However, the data is generated entirely by the under-parameterized TVP model as $\hat{\pi}_t \to \bar{\pi}$, which is close to 1. A correctly specified model fails to generate the data because of feedback.

5. Simulations

We report numerical simulations illustrating the convergence described by Theorem 4.1. We run the model for 10,000 periods and record $\pi_{10,000}$. A typical distribution of $\pi_{10,000}$ over the sample paths is bimodal, concentrating at either 1 or 0. We plot the results by taking the average of $\pi_{10,000}$ over 160 sample paths. We fix $\sigma = 0.004$ and $\rho = .99$. To control the strength of feedback, we change the value of α. We experiment with two values of $\alpha = 0.96$ and $\alpha = 0.25$, representing strong and weak (expectational) feedback. Since the analytic result obtains by taking $\sigma_v \to 0$ and $\sigma_f \to 0$, we choose small values for σ_v and σ_f.

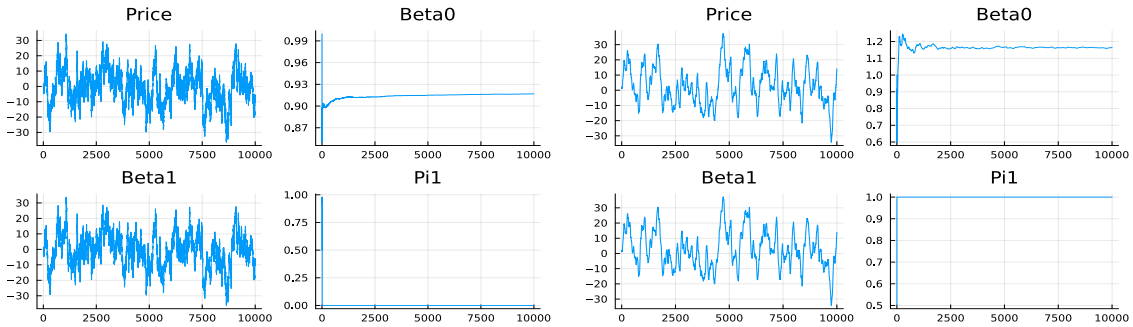


Figure 1: Two examples of sample paths: $\pi_t \to 0$ in the left panel, and $\pi_t \to 1$ in the right panel. We observe both sample paths for a given parameter value, but the proportion of the kind of sample paths in the right panel increases as $(\sigma_v, sigma_f) \to 0$.

In the left panel, $\pi_t \to 0$, while in the right panel $\pi_t \to 1$. By the end of the simulations, the distribution of π_t is bimodal, (mostly) concentrating at 0 or 1. The probability assigned to the TVP model becomes larger and converges to 1, as $\sigma_f, \sigma_v \to 0$. Figure 5 shows the average values of $\pi_{10,000}$ over 160 sample paths for each combination of (σ_v, σ_f). In the left panel where $(\alpha, \rho) = (0.96, 0.99)$ so that $\alpha\rho > 0.5$. As Theorem 4.1 claims, $\pi_{10,000}$ converges to 1. In the right panel where $(\alpha, \rho) = (0.25, 0.99)$, $\pi_{10,000}$ remains close to 0.

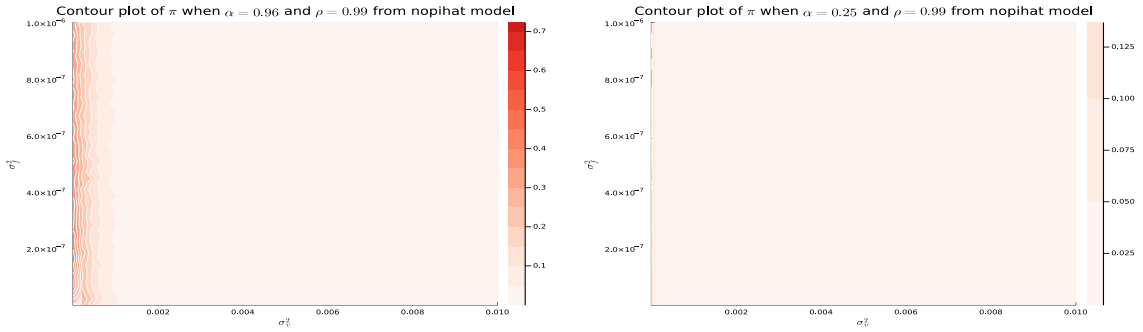


Figure 2: Each point in the panel shows the value of $\pi_{10,000}$ for a given value of (σ_v, σ_f). The left panel shows that with a strong feedback $(\alpha\rho > 0.5)$, $\pi_t \to 1$ as $(\sigma_v, \sigma_f) \to 0$. The right panel shows that if the feedback is weak $(\alpha\rho < 0.5)$, $\pi_t \to 0$, even if $(\sigma_v, \sigma_f) \to 0$.

We repeat the same simulations for the augmented CP model and the TVP model. Since the augmented CP model accurately captures the functional form of the aggregation process by the policymaker, the augmented CP model fares better than the CP model that ignores the policymaker's aggregation process. Theorem B.3 shows the convergence of π_t as $(\sigma_b, \sigma_f) \to 0$. To confirm numerically Theorem B.3, we need to run many extremely long simulations for small values of (σ_b, σ_f), since the evolution of π_t slows down significantly as $(\sigma_b, \sigma_f) \to 0$. Still, for the small values of (σ_b, σ_f) with moderately long simulations of 10,000 periods, we can observe that $\pi_{10,000}$ forms the binomial distribution. As expected,

the probability mass of $\pi_{10,000}$ moves from 0 to 1 as $(\sigma_b, \sigma_f) \to 0$. Interestingly, the probability mass around 1 moves slightly downward since $\pi_t \to \overline{\pi} < 1$.

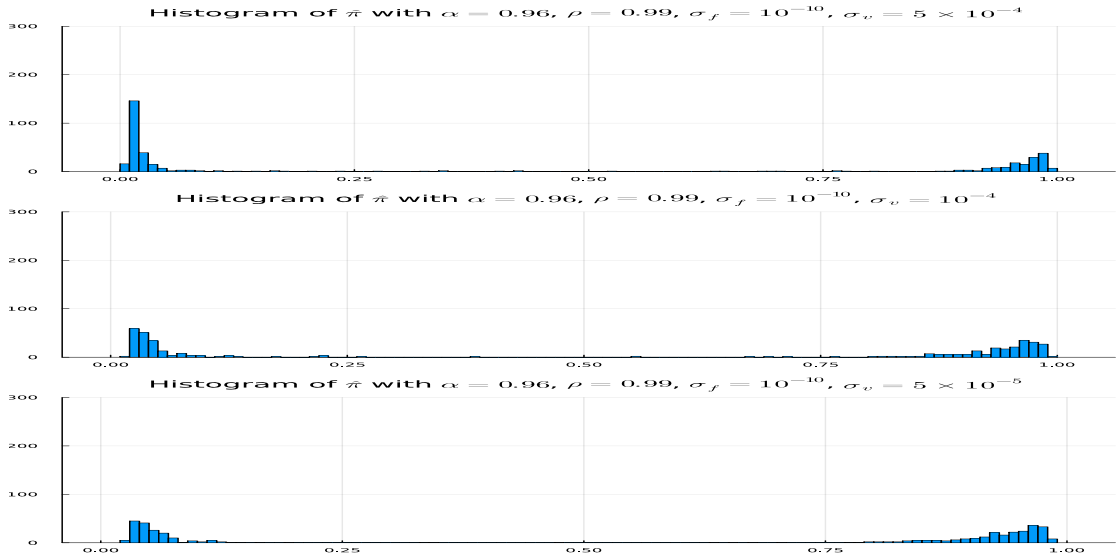


Figure 3: The probabililty mass is concentrated at 0 or 1 as $t \to \infty$. As $(\sigma_v, \sigma_f) \to 0$, the limit probability mass at 0 vanishes.

We report the two sample paths of variables of interest. Since $\pi_t \to \overline{\pi} < 1$, the sample path of π_t entails significant volatility. We plot the sample path of $\hat{\pi}_t$ instead. Being the sample average of π_t, the sample path of $\hat{\pi}_t$ better shows the trend of π_t.

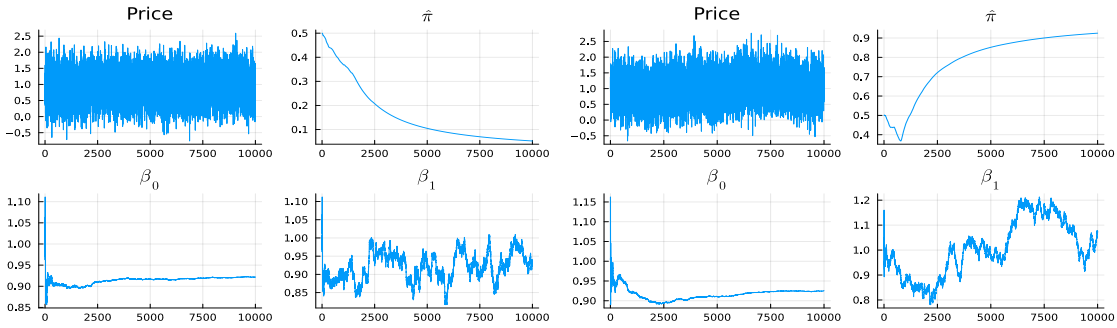


Figure 4: We observe both types of sample paths for a given set of parameter values. As $(\sigma_v, \sigma_f) \to 0$, the second type of sample path becomes dominant.

As (σ_v, σ_f) decreases, the proportion of the sample paths of π_t converging to $\overline{\pi}$ increases toward 1. Because of the volatility of π_t and the slow convergence speed for small (σ_v, σ_f), we observe that moderately large value of average $\pi_{10,000}$ for a combination of small values of (σ_v, σ_f). Since the augmented CP model is correctly specified, $\pi_t \to 0$ if the feedback

is weak or moderate. With the CP model, the TVP model can dominate if $\alpha\rho > 0.5$. The TVP model needs $\alpha\rho$ significantly higher than 0.5 with the augmented CP model. Theorem B.3 shows that the average value of π_t is bounded away from 0 if $\alpha\rho$ is sufficiently large. Even if the augmented CP model is correctly specified, it generates the data with a probability strictly less than 1 when the feedback is strong.

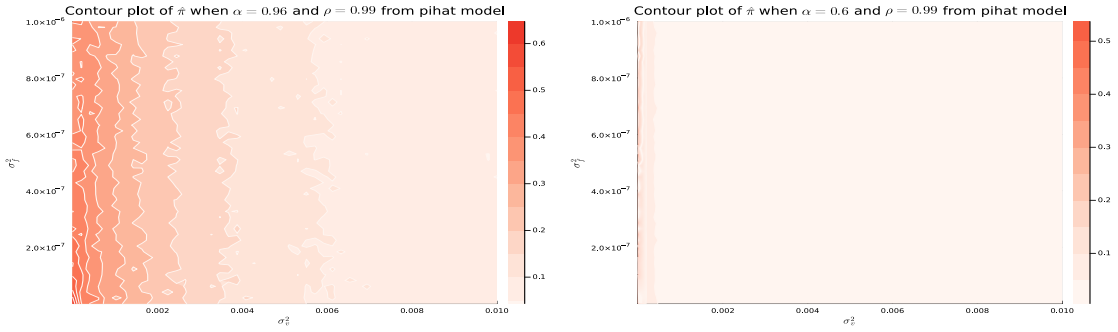


Figure 5: If the feedback is very strong, the expected value of π_t converges to $\overline{\pi} > 0$ as $(\sigma_v, \sigma_f) \to 0$ as shown in the left panel. If the back is not strong, $\pi_t \to 0$ as $(\sigma_v, \sigma_f) \to 0$ as shown in the right panel.

6. Concluding Remarks

Parameter instability is a fact of life for applied econometricians. One explanation is that the economy is exogenously nonstationary. If a model fails to recognize the evolving nature of the economy, its parameters will be unstable. Parameter instability arises from model misspecification. Another commonly advanced explanation is that parameter instability reflects a more subtle form of model misspecification, which occurs when models fail to account for the Lucas Critique. Even if the background environment is time-invariant, if a model misspecifies expectational feedback, its parameters will become unstable. However, by itself, misspecified feedback cannot explain persistent parameter drift. As Sargent (1999, p. 16) notes,

> *Observations are drawn from rational expectations equilibria...would not provide significant evidence of parameter drift, even for misspecified models...Interpreting evidence of coefficient drift as evidence of model misspecification requires an alternative model that somehow causes the moments used in constructing estimators to depend on calendar time....*

Our paper has provided one such model. When misspecified feedback is combined with *competition* among models, a competition that takes the form of Bayesian model averaging, then simple time-varying parameters models can become self-confirming. This occurs even when the TVP model is under-parameterized and when the CP alternative model is quite sophisticated in that it is aware of the nature of the competition. In time-invariant self-referential environments, otherwise, correctly specified models have difficulty competing against more adaptive models that feature time-varying parameters, even when the TVP models are misspecified. As a result, the economy can become *endogenously* nonstationary.

These results raise the obvious question of what can be done to avoid them. The suboptimal self-confirming equilibrium here does not arise from insufficient off-equilibrium-path experimentation, as in Fudenberg and Levine (2009). The decisionmaker gives the CP model every opportunity to succeed. In fact, he gives it some important advantages. Instead, the root cause of the problem is a failure to understand the nature of feedback and endogeneity in the economy. In this era of Big Data and Machine Learning, one sometimes hears the argument that theory has become passe. Our paper provides a warning against such arguments. Although a first-year economics student who proposes a Rational Expectations equilibrium model would lose badly against his Computer Science counterpart, who is an expert at fitting data in real-time, a second-year economics student who has learned the lessons of the Lucas Critique would know enough not to play the game. He would realize that the rules are stacked against him. Instead, he would ask to change the game, e.g., by comparing cross-regime outcomes.

Appendix A. Proof of Proposition 4.1

The proof follows the same reasoning as in Cho and Kasa (2017). We cut and pasted the proof in Cho and Kasa (2017) to make the paper self-contained because we use some important concepts in Appendix B.

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

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

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

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

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

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

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

To compare the speed of evolution, we calculate

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

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

Given $\sigma_v > 0$,

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

remains finite with probability 1. On the other hand,

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

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

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

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

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

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

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

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

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

A.2. **Dynamics of π_t.** To study the dynamics of π_t it is useful to rewrite (B.43) as follows

$$\pi_{t+1} = \pi_t + \pi_t(1-\pi_t)\left[\frac{A_{t+1}(1)/A_{t+1}(0) - 1}{1 + \pi_t(A_{t+1}(1)/A_{t+1}(0) - 1)}\right] \tag{A.28}$$

which has the familiar form of a discrete-time replicator equation, with a stochastic, state-dependent fitness function determined by the likelihood ratio. Equation (A.28) reveals a lot about the model averaging dynamics. First, it is clear that the boundary points $\pi = \{0, 1\}$ are trivially stable fixed points, since they are absorbing. Second, we can also see that there could be an interior fixed point, where $\mathsf{E}(A_{t+1}(1)/A_{t+1}(0)) = 1$. However, we shall also see that this fixed point is unstable. So we know already that π_t will spend most of its time near the boundary points.

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

Proof. With two full-support probability distributions, we can never conclude that a history of any finite length could not have come from either of the distributions. Slightly more formally, if the distributions have full support, they are mutually absolutely continuous, so the likelihood ratio in eq. (A.28) is strictly bounded between 0 and some upper bound B. To see why $\pi_t < 1$ for all t, notice that $\pi_{t+1} < \pi_t + \pi_t(1 - \pi_t)M$ for some $M < 1$, since the likelihood ratio is bounded by B. Therefore, since $\pi + \pi(1-\pi) \in [0, 1]$ for $\pi \in [0, 1]$, we have

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

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

Since the distributions here are assumed to be Gaussian, they obviously have full support, so Proposition A.1 applies. Although the boundary points are unattainable in finite time, the replicator equation for π_t in (A.28) makes it clear that π_t will spend most of its time near these boundary points, since the relationship between π_t and π_{t+1} has the familiar logit function shape, which flattens out near the boundaries. As a result, π_t evolves very slowly near the boundary points. In fact, we shall now show that it evolves even more slowly than the t^{-1} time-scale of $\beta_t(0)$. This means that when studying the dynamics of the coefficient estimates near the boundaries, we can treat π_t as fixed.

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

Lemma A.2.

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

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

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

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

since $\{\pi_t\} \in \Pi_0$. Depending upon the rate of convergence (or the time scale according to which π_t converges to π^*), we have to treat π_t has already converged to π^*.[4]

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

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

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

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

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

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

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

Since

$$p_t = \gamma + \alpha\left([(1-\pi_t)\beta_t(0) + \pi_t\beta_t(1)] + \delta\frac{(1-\pi_t)\rho}{1-\alpha\rho}f_t\right) + \delta f_t + \sigma\epsilon_t,$$

the forecast error of model 1 is

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

Since $\beta_t(1)$ evolves according to (B.38) and $\mathsf{E}(f) = 0$,

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

Define

$$\overline{\beta}(1) = \frac{\gamma + \alpha(1-\pi_t)\beta_t(0)}{1-\alpha\pi_t} \tag{A.29}$$

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

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

or equivalently,

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

where $\mathsf{E}\xi_t^2 \sim \sigma_v^2$ and $\mathsf{E}\xi_t \to 0$.

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

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

by substituting $\overline{\beta}(1)$. Then, we can write the forecast error of $\mathcal{M}_0$ as

$$\begin{aligned} & p_t - p_t^e(0) \\ = \;& \gamma + \alpha\left[\pi_t\left(\frac{\gamma + \alpha(1-\pi_t)\beta_t(0)}{1-\alpha\pi_t} + \xi_t\right) + (1-\pi_t)\left(\beta_t(0) + \frac{\delta\rho f_t}{1-\alpha\rho}\right)\right] + \delta f_t + \sigma\epsilon_t - \left[\beta_t(0) + \frac{\delta}{1-\alpha\rho}f_t\right] \\ = \;& \frac{\gamma - (1-\alpha)\beta_t(0)}{1-\alpha\pi_t} + \alpha\pi_t\xi_t - \delta\frac{\alpha\rho\pi_t}{1-\alpha\rho}f_t + \sigma\epsilon_t. \end{aligned}$$

□

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-known results from stochastic approximation, we know that the asymptotic properties of ϕ_t are determined by the stability properties of the following ordinary differential equation (ODE)

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

which has a unique stable point

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

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

A.3. **Time scale of π_t.** Given any $\alpha \geq 1$, a simple calculation shows

$$t^\alpha(\pi_t - \pi_{t-1}) = \frac{t^\alpha(e^{(t-1)\phi_{t-1}} - e^{t\phi_t})}{(1+e^{t\phi_t})(1+e^{(t-1)\phi_{t-1}})}.$$

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

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

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

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

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

A.4. **Proof of the Theorem.** It is helpful to summarize our findings on the time scale of three stochastic processes: π_t, $\beta_t(0)$ and $\beta_t(1)$. As indicated by (A.28), π_t evolves quickly in the interior of $[0,1]$. However, no sample path converges to $\pi^* \in (0,1)$ with a positive probability. After π_t enters a small neighborhood of $\{0,1\}$, the evolution of π_t slows down significantly. Around the neighborhood of $\{0,1\}$, we have a hierarchy of time scale among three stochastic processes. $\beta_t(1)$ evolves according to a faster time scale than $\beta_t(0)$, which evolves at a faster time scale than π_t.

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

$$\mathcal{S} = \left\{(\pi, \beta(0), \beta(1)) \mid \beta(1) = \overline{\beta}(1)\right\} \tag{A.30}$$

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

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

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

with probability 1, which implies that

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

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

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

where the expectation is taken with respect to the stationary distribution in the limit as $t \to \infty$.

It is convenient to consider the deterministic dynamics in terms of the time scale of $\beta_t(0)$. The domain of attraction for $(\pi, \beta(0), \beta(1)) = \left(0, \frac{\gamma}{1-\alpha}, \frac{\gamma}{1-\alpha}\right)$ is

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

where $A_t(0)$ and $A_t(1)$ are the perceived likelihood functions. The expectation operator is based on the actual probability distribution. Since $\beta(1)$ evolves at a faster time scale,

$$\beta(1) = \frac{\gamma + \alpha(1-\pi_t)\beta(0)}{1-\alpha\pi_t}.$$

After substituting $\beta(1)$,

$$\begin{aligned}
\Psi(\beta(0), \pi_t) &\equiv \mathsf{E} \log \frac{A_t(0)}{A_t(1)} \\
&= -\frac{1}{2(\Sigma_t(0)+\sigma^2)} \left[\alpha^2 \pi_t^2 \sigma_{\xi_t}^2 + \sigma^2 + \left[\frac{\delta\alpha\rho\pi_t}{1-\alpha\rho} \right]^2 \frac{\sigma_f^2}{1-\rho^2} + \left(\frac{\gamma}{1-\alpha} - \beta(0) \right)^2 \left(\frac{\gamma}{1-\alpha\pi_t} \right)^2 \right] \\
&\quad + \frac{1}{2(\Sigma_t(1)+\sigma^2)} \left[(1-\alpha\pi_t)^2 \sigma_{\xi_t}^2 + \sigma^2 + \left[\frac{\delta(1-\alpha\rho\pi_t)}{1-\alpha\rho} \right]^2 \frac{\sigma_f^2}{1-\rho^2} \right] \\
&\quad + \frac{1}{2} \log \left[\frac{\Sigma_t(1)+\sigma^2}{\Sigma_t(0)+\sigma^2} \right].
\end{aligned}$$

Let us delineate $\mathcal{D}_0$ over $(\beta(0), \pi)$. Note that $\mathcal{D}_0$ is symmetric around

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

where π_t is peaked. One can easily check that

$$\frac{d\pi_t}{d\beta(0)} \geq 0$$

if and only if $\beta(0) \leq \frac{\gamma}{1-\alpha}$.

As we calculate the probability that $(\beta(0), \pi_t)$ escapes from $\mathcal{D}_0$, we are interested in the largest value of π_t in $\mathcal{D}_0$, and the maximum width of $\mathcal{D}_0$.

A simple calculation shows that if $\pi_t \geq \frac{1}{2\alpha\rho}$ at $\beta(0) = \frac{\gamma}{1-\alpha}$, then $\Psi(\beta(0), \pi_t) > 0$ for a sufficiently small $\sigma_{\xi_t}^2$. Since we let $\sigma_v \to 0$, we conclude that the largest value of π_t in $\mathcal{D}_0$ cannot exceed $\frac{1}{2\alpha\rho}$.

The widest part of $\mathcal{D}_0$ is where $\pi_t = 0$. Since $\Psi(\beta(0), \pi_t)$ is a quadratic function of $\beta(0)$, there are two values of $\beta(0)$ satisfying

$$\Psi(\beta(0), 0) = 0.$$

Let $\overline{\beta} > \underline{\beta}$ be the two solutions of the above equation. We can write

$$\overline{\beta} = \frac{\gamma}{1-\alpha} + \Delta \quad \text{and} \quad \underline{\beta} = \frac{\gamma}{1-\alpha} - \Delta$$

where

$$\Delta = \frac{\delta\sigma_f}{\gamma(1-\alpha\rho)\sqrt{1-\rho^2}}.$$

Note that for given $\Sigma(1) > 0$ and $\sigma_f^2 > 0$,

$$\lim_{\sigma^2 \to 0} \Delta = 0.$$

While $\frac{\partial\Delta}{\partial\sigma_f^2} > 0$, the sign of $\frac{\partial\Delta}{\partial\sigma_v^2}$ is not obvious. It is possible that $\frac{\partial\Delta}{\partial\sigma_v^2} < 0$, as the first term of Δ reduces as $\Sigma(1)$ increases. As a result, if $\frac{\sigma_v^2}{\sigma_f^2}$ is small, Δ becomes small, which helps Δ approaches to 0 for a given value of σ^2.

Note that if $\sigma_v^2 = 0$, $\sigma_{\xi_t}^2$ and $\Sigma_t(1) = 0$, implying that $\Psi(\frac{\gamma}{1-\alpha}, 0) > 0$, which is necessary and sufficient for the existence of $\overline{\beta} > \underline{\beta}$ such that

$$\overline{\beta} > \frac{\gamma}{1-\alpha} > \underline{\beta}. \tag{A.31}$$

We can choose $\sigma^2 > 0$ sufficiently small so that (A.31) holds.

We calculate a lower bound of the escape probability from $\mathcal{D}_0$. Consider the probability of escape through the boundary of $\mathcal{D}_0$ along $\pi_t = 0$. Since the CP model is correctly specified if $\pi_t = 0$, the forecasting error from the CP model is $\sigma\epsilon_t$, whose variance is σ^2.

Let us consider the probability that $(\pi_t, \beta_t(0))$ escape $\mathcal{D}_0$ through the neighborhood of the boundary where $\pi_t = 0$. Since the perturbation is Gaussian, its rate function

$$\Lambda_+^0 = \frac{\delta^2 \sigma_f^2}{\gamma^2 (1-\alpha\rho)^2 (1-\rho^2)\sigma^2}.$$

is the ratio of the square of the radius of the neighborhood and the variance of the forecasting error of the CP model at $p = \frac{\gamma}{1-\alpha}$, which is σ^2. Since the exit point $(\overline{\beta}, 0), (\underline{\beta}, 0)$ on the boundary of $\mathcal{D}_0$ may not be the most likely exit point, the probability of escape from $\mathcal{D}_0$ is through the most likely exit point along the boundary of $\mathcal{D}_0$ is higher than the probability through $(\overline{\beta}, 0)$ or $(\underline{\beta}, 0)$. Let Λ_+^* be the large deviation rate function of the escape probability from $\mathcal{D}_0$. Clearly,

$$0 < \Lambda_+^* \leq \Lambda_+^0 \tag{A.32}$$

Given σ^2, we choose $\sigma_f^2 \to 0$ to make $\Lambda_+^0 \to 0$, implying $\Lambda_+^* \to 0$ making it eaiser to escape from $\mathcal{D}_0$.

Next, we calculate an upper bound of the escape probability from the complement $\mathcal{D}_0^c$ of $\mathcal{D}_0$. Since $\pi_t = 1$ at the stable point outside of $\mathcal{D}_0$, let us consider $\Psi(\beta(1), 1)$ after $t \to \infty$

$$\begin{aligned}
\Psi(\beta(1),1) \quad = \quad & -\frac{1}{2\sigma^2}\left[\alpha^2\sigma_{\xi_t}^2 + \sigma^2 + \left[\frac{\delta\alpha\rho}{1-\alpha\rho}\right]^2 \frac{\sigma_f^2}{1-\rho^2} + \left(\frac{\gamma}{1-\alpha} - \beta_0\right)^2 \left(\frac{\gamma}{1-\alpha}\right)^2\right] \\
& + \frac{1}{2(\Sigma_t(1)+\sigma^2)}\left[(1-\alpha)^2\sigma_{\xi_t}^2 + \sigma^2 + \left[\frac{\delta(1-\alpha\rho)}{1-\alpha\rho}\right]^2 \frac{\sigma_f^2}{1-\rho^2}\right] \\
& + \frac{1}{2}\log\left[1 + \frac{\Sigma_t(1)}{\sigma^2}\right].
\end{aligned}$$

Note that the sum in the first two lines is negative for any values of σ_f and σ^2 because $\alpha > \frac{1}{2}$. Note also that it becomes more negative as $\sigma^2 \to 0$ for a fixed σ_v^2 at the linear rate of $1/\sigma^2$, while the third line is positive and increases at a slower logarithmic rate. Thus, for a fixed σ_v^2 satisfying (A.31), we can choose $\sigma^2 > 0$ sufficiently small so that $\Psi(\beta(1), 1) < 0$.

Recall that

$$\frac{1-\pi_t}{\pi_t} = \prod_{k=1}^{t} \frac{A_t(0)}{A_t(1)}.$$

To escape from the neighborhood of $\pi_t = 1$ into $\mathcal{D}_0$, it is necessary (if not sufficient) that $\pi_t \leq \frac{1}{2\alpha}$, since

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

Let us calculate the probability $\pi_t \leq \frac{1}{2\alpha}$, which is an upper bound of the escape probability from $\pi_t = 1$ into $\mathcal{D}_0$. The event is equivalent to

$$\frac{1-\pi_t}{\pi_t} = \prod_{k=1}^{t} \frac{A_k(0)}{A_k(1)} \geq 2\alpha - 1 > 0$$

or

$$\frac{1}{t}\sum_{k=1}^{t} \log \frac{A_k(0)}{A_k(1)} \geq \frac{1}{t}\log(2\alpha - 1).$$

We know that

$$\frac{1}{t}\sum_{k=1}^{t}\log\frac{A_k(0)}{A_k(1)} \to \mathsf{E}\log\frac{A_k(0)}{A_k(1)} < 0.$$

After normalizing the random variable, we have

$$\frac{1}{t}\sum_{k=1}^{t}\frac{\log\frac{A_k(0)}{A_k(1)} - \mathsf{E}\log\frac{A_k(0)}{A_k(1)}}{\sqrt{\mathsf{var}(\log\frac{A_k(0)}{A_k(1)})}} \equiv \phi_t^*(\sigma_v^2,\sigma^2,\sigma_f^2)$$

where $\mathsf{var}(\cdot)$ is the variance of the random variable. Since $\log\frac{A_k(0)}{A_k(1)}$ is i.i.d., the variance remains the same over k. We are interested in

$$\lim_{t\to\infty}\mathsf{P}\left(\phi_t^*(\sigma_v^2,\sigma^2,\sigma_f^2) > \frac{t^{-1}\log(2\alpha-1) - \mathsf{E}\log\frac{A_t(0)}{A_t(1)}}{\sqrt{\mathsf{var}(\log\frac{A_t(0)}{A_t(1)})}}\right)$$

which is the probability that the sample average $\phi_t^*(\sigma_v^2,\sigma^2,\sigma_f^2)$ of i.i.d. variables with mean zero becomes larger than a threshold. More precisely, we are interested in the rate at which the probability vanishes

$$-\lim_{t\to\infty}\frac{1}{t}\log\mathsf{P}\left(\phi_t^*(\sigma_v^2,\sigma^2,\sigma_f^2) > \frac{t^{-1}\log(2\alpha-1) - \mathsf{E}\log\frac{A_t(0)}{A_t(1)}}{\sqrt{\mathsf{var}(\log\frac{A_t(0)}{A_t(1)})}}\right).$$

The term inside of log is the probability that the $\pi_t < \frac{1}{2\alpha}$. Since $\{\pi_t \geq \frac{1}{2\alpha}\} \subset \mathcal{D}_0^c$, the probability cannot be smaller than the escape probability from $\mathcal{D}_0^c$. Let Λ_-^* be the large deviation rate function for the escape probability from $\mathcal{D}_0^c$. Clearly,

$$\Lambda_-^* \geq -\lim_{t\to\infty}\frac{1}{t}\log\mathsf{P}\left(\phi_t^*(\sigma_v^2,\sigma^2,\sigma_f^2) > \frac{t^{-1}\log(2\alpha-1) - \mathsf{E}\log\frac{A_t(0)}{A_t(1)}}{\sqrt{\mathsf{var}(\log\frac{A_t(0)}{A_t(1)})}}\right). \tag{A.33}$$

We claim that the right hand side of (A.33) is bounded away from 0, as $t\to\infty$. Since we normalize the random varaible, $\phi_t^*(\sigma_v^2,\sigma^2,\sigma_f^2)$ is the sample mean of standard i.i.d. random variables. The right-hand side of the inequality converges to a positive bound, which is an increasing function of $(2\alpha-1)\sigma_\xi^2$ where σ_ξ^2 is selected to satisfy (A.31). Cramer's theorem says that the tail of the sample average vanishes at a positive rate, implying that the right hand side of (A.33) is positive (possibly, $+\infty$).

Let us summarize the results we have obtained. First, by choosing a sufficiently small $\sigma_v > 0$, we can make σ_ξ^2 so small that (A.31) holds to ensure that $\pi_t = 0$ is a locally stable point. Since the radius of the neighborhood of the sample mean of the normalized i.i.d. random variables is bounded away from 0, the rate function $\Lambda_-^* > 0$ of the escape probability is also bounded away from 0.

Finally, choose $\sigma_f^2 > 0$ sufficiently small so that the rate function Λ_+^* converges to 0 so that $\Lambda_+^* < \Lambda_-^*$. The large deviation rate function of the escape probability from $\mathcal{D}_0$ (where the overparameterized model dominates) is smaller than the escape probability from $\mathcal{D}_0^c$ (where the underparameterized model dominates). The escape probability from $\mathcal{D}_0$ is exponentially larger than that from $\pi_t = 1$ into $\mathcal{D}_0$. The duration time of using the underparameterized model is exponentially longer than the duration time of using the overparameterized model.

Appendix B. Proof of Theorem B.3

Let us rewrite the system of recursive equations for reference.

$$\begin{aligned}
\beta_{t+1}(0) &= \beta_t(0) + \left(\frac{\Sigma_t(0)}{\sigma^2 + \Sigma_t(0)}\right)(p_t - p_t(0)) && \text{(B.34)}\\
p_t(0) &= (1-\hat{\pi}_t)\left[\beta_t(0) + \frac{\delta}{1-\alpha\rho} f_t\right] + \hat{\pi}_t \beta_t(1) && \text{(B.35)}\\
\hat{\pi}_t &= \hat{\pi}_{t-1} + \frac{1}{t}(\pi_t - \hat{\pi}_{t-1}) && \text{(B.36)}\\
\Sigma_{t+1}(0) &= \Sigma_t(0) - \frac{(\Sigma_t(0))^2}{\sigma^2 + \Sigma_t(0)} && \text{(B.37)}\\
\beta_{t+1}(1) &= \beta_t(1) + \left(\frac{\Sigma_t(1)}{\sigma^2 + \Sigma_t(1)}\right)(p_t - p_t(1)) && \text{(B.38)}\\
p_t(1) &= \beta_t(1) && \text{(B.39)}\\
\Sigma_{t+1}(1) &= \Sigma_t(1) - \frac{(\Sigma_t(1))^2}{\sigma^2 + \Sigma_t(1)} + \sigma_v^2 && \text{(B.40)}\\
\hat{p}_t &= (1-\pi_t)p_t(0) + \pi_t p_t(1) && \text{(B.41)}\\
p_t &= \gamma + \alpha \mathsf{E}_t \hat{p}_{t+1} + \delta f_t + \sigma \epsilon_t \\
&= \gamma + \alpha(\pi_t + (1-\pi_t)\hat{\pi}_t)\beta_t(1) + \alpha(1-\pi_t)(1-\hat{\pi}_t)\beta_t(0) \\
&\quad + \left(\alpha(1-\pi_t)(1-\hat{\pi}_t)\frac{\delta\rho}{1-\alpha\rho} + \delta\right) f_t + \epsilon_t. && \text{(B.42)}\\
\frac{1}{\pi_{t+1}} - 1 &= \frac{A_{t+1}(0)}{A_{t+1}(1)}\left(\frac{1}{\pi_t} - 1\right) && \text{(B.43)}
\end{aligned}$$

where $A_t(i)$ is the "perceived" likelihood function of the augmented CP model if $(i = 0)$ and the TVP model if $(i = 1)$:

$$A_t(i) = \frac{1}{\sqrt{2\pi(\sigma^2 + \Sigma_t(i))}} \exp\left[-\frac{(p_t - p_t(i))^2}{2\pi(\sigma^2 + \Sigma_t(i))}\right] \tag{B.44}$$

$\forall i, t$.

Since $\beta_t(1)$ evolves at a faster time scale than $\beta_t(0)$, we fix $\beta_t(0)$ to calculate the limit of the mean of $\beta_t(1)$ that moves according to (B.38). Given that the decision-maker aggregates $p_t(1)$ and $p_t(0)$ according to (B.41), (B.42) becomes

$$\begin{aligned}
p_t &= \gamma + \alpha(\pi_t + (1-\pi_t)\hat{\pi}_t)\beta_t(1) + \alpha(1-\pi_t)(1-\hat{\pi}_t)\beta_t(0) \\
&\quad + \left(\alpha(1-\pi_t)(1-\hat{\pi}_t)\frac{\delta\rho}{1-\alpha\rho} + \delta\right) f_t + \epsilon_t. && \text{(B.45)}
\end{aligned}$$

For a fixed $\beta_t(0) = \beta(0)$, $\beta_t(1)$ converges to satisfy

$$\mathsf{E}\left[p_t - p_t(1)\right] = 0$$

as $t \to \infty$ so that

$$\beta_t(1) = \frac{\gamma + \alpha(1-\pi_t)(1-\hat{\pi}_t)\beta_t(0)}{1 - \alpha(\pi_t + (1-\pi_t)\hat{\pi}_t)} + \xi_t \tag{B.46}$$

implying that

$$\beta_t(1) - \beta_t(0) = \frac{\gamma - (1-\alpha)\beta_t(0)}{1 - \alpha(\pi_t + (1-\pi_t)\hat{\pi}_t)} + \xi_t \tag{B.47}$$

where $\sigma_\xi^2 \to 0$ as $\sigma_v \to 0$. Thus, the variance of the forecasting error of $p_t(1)$ converges to

$$\mathsf{E}(p_t - p_t(1))^2 = A\sigma_\xi^2 + \sigma^2\left(\alpha(\pi_t + (1-\pi_t)\hat{\pi}_t)\right)^2 \Sigma(1) + \left((1-\pi_t)(1-\hat{\pi}_t)\frac{\alpha\delta\rho}{1-\alpha\rho} + \delta\right)^2 \sigma_f^2 \tag{B.48}$$

for some $A < 0$. After substituting (B.46) into (B.45), we have

$$p_t = \frac{\gamma + \alpha(1-\pi_t)^2\beta_t(0)}{1 - \alpha(\pi_t + (1-\pi_t)\pi_t)} + \left(\alpha(1-\pi_t)^2\frac{\delta\rho}{1-\alpha\rho} + \delta\right) f_t + \epsilon_t. \tag{B.49}$$

$\beta_t(0)$ converges to a value satisfying

$$\mathsf{E}p_t - \beta_t(0) = 0$$

implying that

$$\beta_t(0) \to \frac{\gamma}{1-\alpha}$$

with probability 1. A simple calculation shows that

$$\beta_t(1) \to \frac{\gamma}{1-\alpha}$$

weakly. As $t \to \infty$, $\Sigma_t(0) \to 0$. Thus, after setting $\Sigma_t(0) = 0$, we can write

$$\begin{aligned}\mathsf{E}(p_t - p_t(0))^2 &= \left(\frac{(1-\alpha)(1-\hat{\pi}_t)}{1-\alpha(\pi_t+(1-\pi_t)\hat{\pi}_t)}\right)^2\left(\frac{\gamma}{1-\alpha}-\beta_t(0)\right)^2 \\ &\quad + \left((1-\pi_t)(1-\hat{\pi}_t)\frac{\alpha\rho\delta}{1-\alpha\rho} - (1-\hat{\pi}_t)\frac{\delta}{1-\alpha\rho}+\delta\right)^2\frac{\sigma_f^2}{1-\delta^2}+\sigma^2+B\sigma_\xi^2 \quad \text{(B.50)}\end{aligned}$$

for some $B > 0$. Since π_t evolves according to (B.43), we can write

$$\phi_t = \phi_{t-1} + \frac{1}{t}\left(\log\frac{A_t(0)}{A_t(1)} - \phi_{t-1}\right)$$

where

$$\phi_t = \frac{1}{t}\log\left[\frac{1-\pi_t}{\pi_t}\right].$$

The associate ODE is

$$\dot{\phi} = \mathsf{E}\log\frac{A_t(0)}{A_t(1)} - \phi.$$

Note that if $\mathsf{E}\log\frac{A_t(0)}{A_t(1)} > 0$, $\pi_t \to 0$, and if $\mathsf{E}\log\frac{A_t(0)}{A_t(1)} < 0$, $\pi_t \to 1$. We identify a condition that $\forall\sigma^2 > 0$,

$$\lim_{\sigma_f\to 0}\lim_{\sigma_v\to 0}\lim_{t\to\infty}\mathsf{E}\log\frac{A_t(0)}{A_t(1)} > 0$$

which is a critical step to analyze the limit points of $(\hat{\pi}_t, \pi_t)$. Note

$$\mathsf{E}\log\frac{A_t(0)}{A_t(1)} = \frac{\mathsf{E}(p_t - p_t(0))^2}{\Sigma_t(0)+\sigma^2} - \frac{\mathsf{E}(p_t - p_t(1))^2}{\Sigma_t(1)+\sigma^2} + \frac{1}{2}\log\frac{\Sigma_t(1)+\sigma^2}{\Sigma_t(0)+\sigma^2}. \quad \text{(B.51)}$$

Since $t \to \infty$ and $\Sigma_t(0) \to 0$, we set $\Sigma_t(0) = 0$ to simplify the expression of (B.51). Then, we let $\sigma_v \to 0$ so that $\Sigma_t(1), \sigma_\xi^2 \to 0$ to further simply the expression. We can show that (B.51) is positive if and only if

$$\begin{aligned}&-\left[\left(1-(1-\pi_t)(1-\hat{\pi}_t)\left(\frac{\alpha\rho\delta}{1-\alpha\rho}-(1-\hat{\pi}_t)\frac{\delta}{1-\alpha\rho}+\delta\right)\right)^2\frac{\sigma_f^2}{1-\delta^2}+\sigma^2\right. \\ &\left.+\left(\frac{(1-\alpha)(1-\hat{\pi}_t)}{1-\alpha(\pi_t+(1-\pi_t)\hat{\pi}_t)}\right)^2\left(\frac{\gamma}{1-\alpha}-\beta_t(0)\right)^2\right)\Bigg] + \left[\left((1-\pi_t)(1-\hat{\pi}_t)\frac{\alpha\delta\rho}{1-\alpha\rho}+\delta\right)^2\frac{\sigma_f^2}{1-\delta^2}+\sigma^2\right].\end{aligned}$$

Note

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

if and only if

$$\begin{aligned}&\left(\frac{(1-\alpha)(1-\hat{\pi}_t)}{1-\alpha(\pi_t+(1-\pi_t)\hat{\pi}_t)}\right)^2\left(\beta_t(0)-\frac{\gamma}{1-\alpha}\right)^2 \\ \leq\ &\frac{\sigma_f^2}{1-\delta^2}\left[\left((1-\hat{\pi}_t)(1-\pi_t)\frac{\alpha\rho\delta}{1-\alpha\rho}+\delta\right)^2 - \left((1-\pi_t)(1-\hat{\pi}_t)\left(\frac{\alpha\rho\delta}{1-\alpha\rho}-(1-\hat{\pi}_t)\frac{\delta}{1-\alpha\rho}+\delta\right)\right)^2\right] \\ =\ &\frac{\sigma_f^2}{1-\delta^2}\left[(1-\hat{\pi}_t)\frac{\delta}{1-\alpha\rho}\right]\left[2(1-\pi_t)(1-\hat{\pi}_t)\left(\frac{\alpha\rho\delta}{1-\alpha\rho}+\delta\right)-(1-\hat{\pi}_t)\frac{\delta}{1-\alpha\rho}\right].\end{aligned}$$

Since the left-hand side of the inequality is positive, the right-hand side must also be positive, implying that the term inside the last bracket on the right-hand side of the inequality must be positive. The term in the last bracket is positive if and only if

$$\alpha\rho < \frac{1+\hat{\pi}_t}{2(1-(1-\hat{\pi}_t)(1-\pi_t))}. \tag{B.52}$$

B.1. **Adaptive Learning.** To analyze the dynamic system with the augmented CP forecaster, we need to consider $\hat{\pi}_t$ in addition to $(\beta_t(0), \pi_t)$, while assuming that $\beta_t(1)$ evolves according to (B.47). Note that $\beta_t(0)$ and $\hat{\pi}_t$ evolves according to the time scale of the sample average. Also, note that π_t evolves at a faster rate than $\beta_t(0)$ and $\hat{\pi}_t$ in the interior of $[0,1]$, but in a small neighborhood of the boundary of $[0,1]$, π_t evolves according to a slower time scale than the other two variables.

Consider $(\beta_t(0), \hat{\pi}_t, \pi_t) \in \mathbb{R} \times [0,1]^2$. Define

$$\mathcal{R}_+ = \left\{ (\beta(0), \hat{\pi}, \pi) \mid \mathsf{E} \log \frac{A(0)}{A(1)} \geq 0 \right\} \tag{B.53}$$

as the area in $\mathbb{R} \times [0,1]^2$ where $\pi_t \to 0$ very quickly. Similarly, define

$$\mathcal{R}_- = \left\{ (\beta(0), \hat{\pi}, \pi) \mid \mathsf{E} \log \frac{A(0)}{A(1)} < 0 \right\} \tag{B.54}$$

as the area in $\mathbb{R} \times [0,1]^2$ where $\pi_t \to \{1,0\}$ very quickly for given $(\beta_t(0), \hat{\pi}_t)$.

For a small value of $\alpha\rho < 1$, $(\beta(0), \hat{\pi}, \pi) = (\frac{\gamma}{1-\alpha}, 0, 0)$ is the globally stable point, implying that the decision-maker believes in the augmented CP model with probability 1 if the feedback is not strong. However, if $\alpha\rho$ is sufficiently close to 1, the dynamic system has another stable point where the decision maker assigns probability weight close to 1 to the (under-parameterized) TVP.

To see this, consider a hyperplane

$$\mathsf{H} = \{(\beta(0), \hat{\pi}, \pi) \mid \pi = \hat{\pi}\} \tag{B.55}$$

where the sample average of π_t coincides with the expected value of π_t. Define

$$\mathsf{H}_+ = \{(\beta(0), \hat{\pi}, \pi) \mid \pi \geq \hat{\pi}\}$$

where the expected value of π_t is larger than $\hat{\pi}_t$. It typically happens if $\pi_t = 1$. Similarly, define

$$\mathsf{H}_- = \{(\beta(0), \hat{\pi}, \pi) \mid \pi \leq \hat{\pi}\}$$

as the closed half space above and below H where the expected value of π_t is larger than $\hat{\pi}_t$. It typically happens if $\pi_t = 0$.

Lemma B.1. *$\mathcal{R}_+ \cap \mathsf{H}$ is a connected set if and only if $\alpha\rho < \frac{2+\sqrt{3}}{4} < 1$.*

Proof. Consider the "ridge line" of $\mathcal{R}_+$. The cross section of $\mathcal{R}_+$ at each value of $\hat{\pi}_t$ is peaked if $\beta_t(0) = \frac{\gamma}{1-\alpha}$. Let $\mathcal{R}_+(\hat{\pi})$ be the cross section of $\mathcal{R}_+$ at $\hat{\pi}$. Define

$$\pi^*(\hat{\pi}) = \max_{\pi} \left\{ \left(\frac{\gamma}{1-\alpha}, \hat{\pi}, \pi \right) \in \mathcal{R}_+ \right\}$$

as the largest value of π in cross section $\mathcal{R}_+(\hat{\pi})$. By the "ridge line" of $\mathcal{R}_+$, we mean the curve

$$\left\{ (\beta(0), \hat{\pi}, \pi) \mid \beta(0) = \frac{\gamma}{1-\alpha}, \pi = \pi^*(\hat{\pi}), \hat{\pi} \in [0,1] \right\}. \tag{B.56}$$

One can easily show that $\mathcal{R}_+ \cap \mathsf{H}$ is connected if and only if (B.56) is connected, which is equivalent to

$$2((1-\pi)^2\alpha\rho + 1 - \alpha\rho) - (1-\pi) = 0$$

has no real solution. A simple calculation shows that the quadratic equation has no real-valued solution if and only if

$$\alpha\rho < \frac{2+\sqrt{3}}{4}.$$

If

$$\alpha\rho \geq \frac{2+\sqrt{3}}{4},$$

$$\bar{\pi} = \frac{4\alpha\rho - 1 + \sqrt{1 - 16\alpha\rho(1-\alpha\rho)}}{4\alpha\rho} \tag{B.57}$$

and

$$\underline{\pi} = \frac{4\alpha\rho - 1 - \sqrt{1 - 16\alpha\rho(1-\alpha\rho)}}{4\alpha\rho} \tag{B.58}$$

are the two solutions. □

Proposition B.2. *Suppose that $t \to \infty$ and then $\sigma_v^2 \to 0$. If*

$$\alpha\rho < \frac{2+\sqrt{3}}{4},$$

then

$$(\beta(0), \hat{\pi}, \pi) = \left(\frac{\gamma}{1-\alpha}, 0, 0\right)$$

is the unique stable solution of the mean dynamics of $(\beta_t(0), \hat{\pi}_t, \pi_t)$ so that

$$(\beta_t(0), \hat{\pi}_t, \pi_t) \to \left(\frac{\gamma}{1-\alpha}, 0, 0\right)$$

weakly. If

$$\alpha\rho > \frac{2+\sqrt{3}}{4},$$

so that $\bar{\pi} > \underline{\pi}$, then

$$(\beta(0), \hat{\pi}, \pi) = \left(\frac{\gamma}{1-\alpha}, \bar{\pi}, \bar{\pi}\right)$$

and

$$(\beta(0), \hat{\pi}, \pi) = \left(\frac{\gamma}{1-\alpha}, 0, 0\right)$$

are two locally stable stationary solutions of the mean dynamics of $(\beta_t(0), \hat{\pi}_t, \pi_t)$.

Proof. Observe that π_t evolves at a faster time scale than $\beta_t(0)$ and $\hat{\pi}_t$ in the interior of $[0,1]$. If $(\beta_t(0), \hat{\pi}_t, \pi_t) \in \mathcal{R}_-$, then $\pi_t \to 1$ (or close to 1, because $\pi_t = 1$ is not reachable), and if $(\beta_t(0), \hat{\pi}_t, \pi_t) \in \mathcal{R}_+$, then $\pi_t \to 0$ (or close to 0, because $\pi_t = 0$ is not reachable). From the perspective of $(\beta_t(0), \hat{\pi}_t)$, $\pi_t \to 1$ and $\pi_t \to 0$ occurs almost instantaneously. Thus, the mean dynamics of $\hat{\pi}_t$ can be approximated by

$$\dot{\hat{\pi}} = \begin{cases} 1 - \hat{\pi}_t & \text{if } (\beta_t(0), \hat{\pi}_t, \pi_t) \in \mathcal{R}_-; \\ -\hat{\pi}_t & \text{if } (\beta_t(0), \hat{\pi}_t, \pi_t) \in \mathcal{R}_+. \end{cases}$$

Let us project the gradient vector field to the space of $(\beta_t(0), \hat{\pi}_t)$ where the variables evolve according to the same time scale as the sample average.

Suppose $\alpha\rho < \frac{2+\sqrt{3}}{4}$. Since $\beta_t(0) \to \frac{\gamma}{1-\alpha}$ with probability 1. It suffices to check whether there is a stable point along the hyperplane where $\beta_t(0) = \frac{\gamma}{1-\alpha}$. If $\pi_t = 1$, $\hat{\pi}_t$ cannot be larger than $2\alpha\rho - 1 \in (0,1)$ since $\frac{1}{2} < \alpha\rho < 1$. As soon as $\hat{\pi}_t > 2\alpha\rho - 1$ while $\pi_t = 1$, $\pi_t \to 0$ quickly which pushes $\hat{\pi}_t$ toward 0. Since $\alpha\rho < \frac{2+\sqrt{3}}{4}$, the "ridge line" of $\mathcal{R}_-$ stays above H along $\beta_t(0) = \frac{\gamma}{1-\alpha}$ for any value of $\hat{\pi}_t \in [0,1]$. Thus, $\hat{\pi}_t \to 0$ while $\pi_t = 0$ in the long run. Thus,

$$(\beta(0), \hat{\pi}, \pi) = \left(\frac{\gamma}{1-\alpha}, 0, 0\right)$$

is the only stable point of the mean dynamics of $(\beta_t(0), \hat{\pi}_t, \pi_t)$.

Suppose $\alpha\rho > \frac{2+\sqrt{3}}{4}$. We know that if a stationary point exists, it should be along the intersection of the hyperplane where $\beta(0) = \frac{\gamma}{1-\alpha}$, the boundary of the intersection between H and $\mathcal{R}_-$ and $(\frac{\gamma}{1-\alpha}, 0, 0)$. If $\alpha\rho > \frac{2+\sqrt{3}}{4}$, then $\bar{\pi} > \underline{\pi}$. Thus, $(\frac{\gamma}{1-\alpha}, \underline{\pi}, \underline{\pi})$ and $(\frac{\gamma}{1-\alpha}, \bar{\pi}, \bar{\pi})$ are two stationary points. Note that the line segment connecting $(\frac{\gamma}{1-\alpha}, \underline{\pi}, \underline{\pi})$ and $(\frac{\gamma}{1-\alpha}, \bar{\pi}, \bar{\pi})$ is in $\mathcal{R}_-$, implying that $\pi_t \to 1$ instantaneously from the perspective of $(\beta_t(0), \hat{\pi}_t)$. As a result, along the line segment

$$\dot{\hat{\pi}} = 1 - \hat{\pi} > 0.$$

Similarly, the line segments connecting $(\frac{\gamma}{1-\alpha}, \underline{\pi}, \underline{\pi})$ and $(\frac{\gamma}{1-\alpha}, 0, 0)$ and and connecting $(\frac{\gamma}{1-\alpha}, \bar{\pi}, \bar{\pi})$ and $(\frac{\gamma}{1-\alpha}, 1, 1)$ belong to $\mathcal{R}_+$ so that $\pi_t \to 0$ instantaneously from the perspective of $(\beta_t(0), \hat{\pi}_t)$, implying

$$\dot{\hat{\pi}} = -\hat{\pi} < 0.$$

Thus, $(\frac{\gamma}{1-\alpha}, \underline{\pi}, \underline{\pi})$ is unstable, while $(\frac{\gamma}{1-\alpha}, \bar{\pi}, \bar{\pi})$ is stable. The stability of $(\frac{\gamma}{1-\alpha}, 0, 0)$ follows from the same logic as the proof of the stability of the same stationary point when $\alpha\rho < \frac{2+\sqrt{3}}{4}$. The only exception is that the domain of attraction changes as $\alpha\rho > \frac{2+\sqrt{3}}{4}$.

Recall that if $\alpha\rho > \frac{2+\sqrt{3}}{4}$, the entire state space is the domain of attraction of $(\frac{\gamma}{1-\alpha}, 0, 0)$, which is the only stable stationary point. If $\alpha\rho < \frac{2+\sqrt{3}}{4}$, a simple check of the gradient vector field shows that

$$\{(\beta(0), \hat{\pi}, \pi) \in \mathcal{R}_+ \cap \mathsf{H}_- \mid \hat{\pi} < \underline{\pi}\} \tag{B.59}$$

is the domain of attraction for $(\frac{\gamma}{1-\alpha}, 0, 0)$. Following the same reasoning, we can show that the domain of attraction for $(\frac{\gamma}{1-\alpha}, \bar{\pi}, \bar{\pi})$ is the complement of (B.59). □

$\forall \epsilon > 0$, define

$$T^\epsilon\left(T, \frac{\gamma}{1-\alpha}, \bar{\pi}\right) = \#\left\{t \le T \mid (\beta_t(0), \hat{\pi}_t) \in \mathcal{N}_\epsilon\left(\left(\frac{\gamma}{1-\alpha}, \bar{\pi}\right)\right)\right\}$$

as the number of periods before T rounds when $(\beta_t(0), \hat{\pi}_t)$ is in ϵ neighborhood of $\left(\frac{\gamma}{1-\alpha}, \bar{\pi}\right)$. Note that $\left(\frac{\gamma}{1-\alpha}, \bar{\pi}\right)$ is not a locally stable point, but a projection of the locally stable point $\left(\frac{\gamma}{1-\alpha}, \bar{\pi}, \bar{\pi}\right)$ to the first two components space.

Theorem B.3. *Suppose that* $\alpha\rho > \frac{2+\sqrt{3}}{4}$.

$$\lim_{\sigma_f \to 0} \lim_{\sigma_v \to 0} \lim_{T \to \infty} \mathsf{E} \frac{T^\epsilon\left(T, \frac{\gamma}{1-\alpha}, \bar{\pi}\right)}{T} = 1. \tag{B.60}$$

Proof. We have used the assumption that $T \to \infty$ by focusing on the limit to set $\Sigma(0) = 0$, and then that $\sigma_v \to 0$ to set $\Sigma(1) = 0$. Note that by letting $\sigma_f \to 0$, the domain of attraction of locally stable point $(\frac{\gamma}{1-\alpha}, 0, 0)$ collapses to the hyperplane where $\beta(0) = \frac{\gamma}{1-\alpha}$. Since the variance of ϵ_t remains bounded away from 0, the large deviation rate function of $(\beta_t(0), \hat{\pi}_t, \pi_t)$ converges to 0 as $\sigma_f \to 0$.

On the other hand, the domain of attraction of $(\frac{\gamma}{1-\alpha}, \bar{\pi})$ expands as $\sigma_f \to 0$, which ensures that the large deviation rate function of $(\beta_t(0), \hat{\pi}_t)$ in the neighborhood of $(\frac{\gamma}{1-\alpha}, \bar{\pi})$ remains bounded away from 0 as $\sigma_f \to 0$.

For a sufficiently small $\sigma_f > 0$, the rate function around $(\frac{\gamma}{1-\alpha}, \bar{\pi})$ becomes larger than the rate function around $(\frac{\gamma}{1-\alpha}, 0)$. Consequently, the duration time of $(\beta_t(0), \hat{\pi}_t)$ staying in the neighborhood of $(\frac{\gamma}{1-\alpha}, \bar{\pi})$ becomes exponentially longer than the duration time in the neighborhood of $(\frac{\gamma}{1-\alpha}, 0)$, from which the conclusion follows. □

References

Berk, R. H. (1966): "Limiting Behavior of Posterior Distributions When the Model is Incorrect," *Annals of Mathematical Statistics*, 37, 51–58.

Blinder, A. S. (1998): *Central Banking in Theory and Practice*. MIT Press.

Blume, L. E., and D. Easley (1982): "Learning to Be Rational," *Journal of Economic Theory*, 26, 340–51.

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

Bray, M. M., and D. M. Kreps (1987): "Rational Learning and Rational Expectations," in *Arrow and the Ascent of Modern Economic Theory*, ed. by G. R. Feiwel, pp. 597–625. New York University Press.

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

——— (2017): "Gresham's Law of Model Averaging," *American Economic Review*, pp. 3589–3616.

Doan, T., R. Litterman, and C. Sims (1984): "Forecasting and Conditional Projection Using Realistic Prior Distributions," *Econometric Reviews*, 3, 1–100.

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

——— (2021): "Equilibrium in Misspecified Markov Decision Processes," *Theoretical Economics*, 16, 717–57.

Evans, G. W., and S. Honkapohja (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.

Fudenberg, D., G. Lanzani, and P. Strack (2021): "Limit Points of Endogenous Misspecified Learning," *Econometrica*, 89, 1065–98.

Fudenberg, D., and D. K. Levine (2009): "Self-Confirming Equilibrium and the Lucas Critique," *Journal of Economic Theory*, 144, 2354–71.

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

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.

Nachbar, J. H. (1997): "Prediction, Optimization, and Learning in Repeated Games," *Econometrica*, 65, 275–309.

Nelson, C. R. (1972): "The Predictive Performance of the FRB-MIT-PENN Model of the U.S. Economy," *American Economic Review*, 62, 902–17.

Olea, J., P. Ortoleva, M. Pai, and A. Prat (2022): "Competing Models," *Quarterly Journal of Economics*, 137, 2419–57.

Department of Economics, Emory University, 1602 Fishburne Dr., Atlanta GA 30322 USA
Email address: In-koo.cho@emory.edu
URL: https://sites.google.com/site/inkoocho

Department of Economics, Simon Fraser University, 8888 University Drive, Burnaby, B.C. Canada. V5A 1S6
Email address: kkasa@sfu.ca
URL: http://www.sfu.ca/~kkasa