

LEARNING AND MODEL VALIDATION

IN-KOO CHO AND KENNETH KASA

ABSTRACT. This paper extends the macroeconomic learning literature by allowing agents to test the specification of their models. It studies the following problem. An agent takes actions based on a possibly misspecified model. The agent is *large*, in the sense that his actions influence the model he is trying to learn about. The agent is aware of potential model misspecification and tries to detect it, in real-time, using an econometric specification test. If his model fails the test, he formulates a new model. If his model passes the test, he uses it to implement a policy based on the provisional assumption that the current model is correctly specified, and will not change in the future.

We claim that this testing and model validation process is an accurate description of most macroeconomic policy problems. Unfortunately, the dynamics produced by this process are not well understood. We make progress on this problem by exploiting results from the large deviations literature. Our analysis can be interpreted as providing a selection criterion for self-confirming equilibria, based on their ‘robustness’. Robust self-confirming equilibria survive repeated specification tests, and are characterized by their large deviation rate functions.

JEL Classification Numbers: C120, E590

1. INTRODUCTION

The macroeconomic learning literature can be seen as progressing through a number of stages. Early contributions focused on convergence questions (Bray (1982), Bray and Savin (1986)). Addressing these questions led naturally to issues of stability and selection among multiple Rational Expectations Equilibria (Sargent (1993)). Once these theoretical issues were resolved, attention shifted to empirical applications. Here the nonstationarity of least-squares learning poses difficult econometric problems. A breakthrough came with the work of Sargent (1999), which introduced the idea of ‘constant gain’ learning. These models feature ‘perpetual learning’ and stationary equilibria, and are therefore better suited to econometric analysis. Finally, although much empirical work remains to be done, the most recent phase of the learning literature has turned its attention to normative questions, related to optimal policy design (Bullard and Mitra (2002), Evans and Honkapohja (2003)).

Looking back on this literature, one is struck by the fact that it always makes one important assumption, namely, that agents are somehow endowed with a *given* model. In the early literature this model conformed to the Rational Expectations Equilibrium, and the question was whether coefficient estimates would converge to their Rational Expectations

Date: October, 2008.

We thank Jim Bullard, Lars Hansen, Seppo Honkapohja, Albert Marcet, and Tom Sargent for helpful discussions. Financial support from the National Science Foundation (ECS-0523620, SES-0720592) is gratefully acknowledged. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

counterparts. More recently, researchers have begun to explore the implications of model misspecification (Sargent (1999) and Evans and Honkapohja (2001)). Still, even here, agents are not allowed to question their models, so that any misspecification necessarily goes undetected. This is unfortunate. From the beginning, one of the main objectives of the learning literature has been to treat agents and their modelers more symmetrically. Although allowing agents to revise coefficient estimates takes a step in this direction, one could argue that economists actually spend *most* of their time searching for better models, not refining estimates of a given model.

Our paper is therefore an attempt to take the next natural step in the learning literature, by allowing agents to *test* the specification of their models. During the past two decades there has been substantial progress on testing and comparing misspecified models.¹ A unifying theme in this work is to view models as inducing probability distributions over the data, and to then compare models using information-theoretic measures of the difference between two probability distributions. Most recent work on comparing misspecified models is based on the Kullback-Leibler Information Criterion (KLIC), which measures the relative entropy between two probability distributions. Models with smaller estimated KLICs are preferred. A key result in this literature is that consistent estimates of the difference between two KLICs can be obtained *without* prior knowledge of the true data-generating process. As a result, models can be compared without having to assume that any of them are correctly specified. Basing model selection on relative entropy also has the virtue that it naturally delivers a measure of model complexity, thereby striking a balance between bias and variance, and allowing models with different dimensions to be sensibly compared. Agents in our model exploit these recent developments.

It is important to keep in mind, however, that these econometric methods presume the data-generating process is exogenous. In our model this is not the case. We study learning by a large agent, like a government policy maker, whose actions influence the data-generating process. As a result, the data *respond to* the agent's own estimation and testing efforts. When it comes to inference, this means there will be a difference between the nominal and actual sizes of our agent's test statistics. We assume this discrepancy goes undetected, which can be interpreted as a form of bounded rationality.

We assume agents have a menu of models, called the *model class*, from which they select the best approximation of the unknown true data generating process. Although in principle we could specify the model class and the testing procedure non-parametrically, this paper opts for the following parametric, information-theoretic approach:² (1) An agent is endowed with a collection of potentially misspecified linear models, each with a finite number of unknown parameters, (2) Each period the agent tests the specification of his current model, (3) If the current model survives the test, the model is updated and used to formulate a policy function, and (4) If the model is rejected, a new model is selected, based on the model with the minimum estimated KLIC statistic. This combined process of estimation, testing, and selection will be called *model validation*.

¹A highly selected sample includes: White (1982), Vuong (1989), Hansen and Sargent (1993), Sin and White (1996), and Rivers and Vuong (2002). White (1994) and Burnham and Anderson (2002) contain textbook treatments.

²It would be an interesting (and ambitious) exercise to combine the nonparametric learning approach of Chen and White (1998) with recent results in nonparametric statistics.

Note that in this validation process a model is only replaced after it fails a specification test. An alternative approach would be to continuously *compare* models. In this case, it might be desirable to switch models before the current reference model actually fails a specification test, if there appears to be a superior alternative. We examine this approach in Cho and Kasa (2008). The dynamics of recursive model comparison are more subtle than those of recursive specification testing, because one needs to continuously monitor the performance of each model, even when it is not being used. We regard both approaches as useful. Whether agents employ specification testing or model comparison likely hinges on case-specific factors, such as the cost of switching models.

Traditional analysis of recursive learning exploits a time-scale separation between the evolution of the data, which operate on a ‘fast’ time-scale, and the evolution of a model’s coefficient estimates, which operate on a ‘slow’ time-scale. Introducing specification testing adds a *third* time-scale, pertaining to the frequency of model switching. Model switches are very rare here, because they occur from self-confirming equilibria, and are therefore ‘large deviation’ events. The fact that each model’s coefficients can be adapted to fit the data it generates is crucial to this result, and it illustrates a key difference between testing with endogenous data and testing with exogenous data.

We show that model validation dynamics can be approximated by a set of Markov-modulated ordinary differential equations. In the limit, as the update gain parameter converges to zero, the agent settles on a *single* model. We can identify this model from its large deviations rate function. In a sense, therefore, our analysis provides a selection criterion for recursive learning models, and can also be interpreted as a refinement of the concept of self-confirming equilibria.

There has been some prior work attempting to link learning with model validation. First, the early work of Bray and Savin (1986) touched on this issue. They ask whether agents could use standard diagnostics, like Chow tests and Durbin-Watson statistics, to detect the time variation in parameters that their own learning behavior generates. Bray and Savin (1986) found that when convergence is slow, agents are generally able to detect the misspecification of their models. In a repeated game context, Foster and Young (2003) allow players to construct, test, and revise simple models of their opponent’s behavior. Hypothesis testing produces convergence to a Nash equilibrium in a relatively strong sense, although testing errors produce rare but recurrent experimentation phases. Perhaps closest in spirit to our analysis is recent work by Branch and Evans (2007). They also study a situation where agents not only update the coefficients of a given model, but also select among alternative parametric models based on their recent forecasting performance. However, they do not base model selection on statistical foundations.

We should clarify at the outset that we make no normative claims in this paper. Instead, we adopt an unabashedly descriptive, or ‘behavioral’, approach to model validation. Our assumptions are motivated by the fact that they seem to be reasonably consistent with the actual policy making process. A prime example concerns the issue of model selection versus model averaging. Standard (Bayesian) decision theory recommends model averaging. However, actual policy making seems more consistent with model selection. Why this is true is an interesting question, but we do not address it in this paper.

The remainder of the paper is organized as follows. Section 2 introduces a few examples of the importance of specification testing in learning models. Section 3 maps our model

validation approach into a standard Stochastic Recursive Algorithm. Section 4 applies results from the large deviations literature to characterize model validation dynamics. Section 5 contains a few concluding remarks, and an appendix contains proofs of some technical results.

2. EXAMPLES

To set the stage, we first present some informal examples designed to illustrate why we think it is important to incorporate specification testing and model validation procedures into the macroeconomic learning literature. We consider two sets of examples, one based on Sargent’s (1999) Phillips Curve model, and the other based on the cobweb model of Branch and Evans (2006).

2.1. Sargent’s Conquest Model. Sargent studied the problem of a Central Bank that wants to minimize a quadratic loss function in inflation and unemployment, but is unsure about the true model. The Bank posits a reduced form regression model of the form,

$$u_n = \beta_0 + \beta_1 \pi_n + \epsilon_n$$

and then tries to learn about this relationship by adaptively updating its coefficients. To account for potential drift in the underlying relationship, the Central Bank discounts past data when updating its estimates. Simulations of this model produce the following time paths of inflation. (Here, and in what follows, we use the same parameter values as in Sargent). The striking feature here is the recurring cycle of gradually rising inflation, and




FIGURE 1. Simulated Inflation Path in Sargent’s Conquest Model

then occasional sharp inflation stabilizations. As noted by Cho, Williams, and Sargent (2002) (CWS), this cycle represents the interplay between the model’s *mean dynamics* and its so-called *escape dynamics*. The mean dynamics reflect the Central Bank’s efforts to eliminate systematic forecast errors. These errors are eliminated once inflation reaches its Nash/Self-Confirming Equilibrium (SCE) value of 5%. The escape dynamics are more exotic. In Sargent’s model, the true data-generating process is a natural rate model, containing no exploitable Phillips Curve trade-off. At the SCE, the Central Bank’s beliefs are free to wander in any direction, and when sequences of positively correlated inflation

and Phillips Curve shocks occur, they cause the Central Bank to revise downward its Phillips Curve slope estimate, and therefore, its inflation target. Since in truth there is no tradeoff, these inflation reductions produce further downward revisions, and the process feeds on itself until inflation reaches the Ramsey outcome of zero inflation. From this point, with no further changes in the inflation target, the Central Bank begins to rediscover the Phillips Curve, due to the presence of inflation shocks acting within the model's natural rate structure. This produces a gradual pull back up to the Nash inflation outcome.

2.1.1. *Sequential t-testing in Sargent's Model.* A natural question at this point is - To what extent is the Central Bank really learning anything here? True, it's revising estimates of a model in light of new data, but in practice policy makers spend most of their time looking for new and improved models, not refining estimates of a *given* model. In Sargent's analysis, the Central Bank never really evaluates the Phillips Curve as theoretical model of inflation and unemployment; it simply reconsiders the strength of an unquestioned trade-off. What if the Central Bank engages in a traditional process of hypothesis testing and model selection? In particular, suppose the Bank entertains the *possibility* that there is no trade-off, perhaps because someone in the research department read a recent presidential address by Milton Friedman. In response, the Bank decides to sequentially test the hypothesis that $\beta_1 = 0$, and if the hypothesis is not rejected, they switch to a vertical Phillips Curve model and set the inflation target to zero. Looking at the *t*-statistics reported in the bottom panel of Figure 1 suggests that this might produce a different result. Once inflation is stabilized, the Bank would clearly switch to a vertical Phillips Curve. In fact, the actual outcome is reported in Figure 2.




FIGURE 2. Sequential *t*-tests in Sargent's Conquest Model

Clearly, allowing the Bank to test hypotheses on the slope coefficient makes no difference, aside from a slight delay in the return to the SCE.³ The fact is, there *is* a correlation between inflation and unemployment, albeit not an exploitable one, and this correlation will lead the Bank to quickly reject the null hypothesis $\beta_1 = 0$, even after it has switched

³Leave aside for the moment the naivete associated with repeatedly applying a *t*-test with a fixed threshold. This will be addressed in the following examples, and makes little difference to the point being made here.

to the vertical Phillips Curve model. The problem, of course, is that the Bank’s model is subject to a fundamental misspecification, based on a misinterpretation of the role of the private sector’s expectations in the inflation process. To break out of its inflation cycle, the Central Bank must engage in a more sophisticated form of specification testing.

2.1.2. *Expected vs. Unexpected Inflation.* Now suppose a young, newly minted PhD arrives in the research department, and having read recent papers by Lucas, Sargent, and Barro, convinces the Bank to try to distinguish between expected and unexpected inflation when estimating its model. That is, it fits an ‘expectations-augmented Phillips Curve’. It turns out, this is quite simple to do in Sargent’s *Conquest* model, since the private sector’s expectations are based directly on the Bank’s own inflation target. Hence, rather than regress unemployment on realized inflation, the Bank just needs to run a multiple regression on its target and the ex post realized inflation shock. Assume the Bank has been conducting business-as-usual until mid-sample, observation 2500, when it suddenly decides to give the new researcher’s ideas a try. The outcome is depicted in Figure 3




FIGURE 3. Expected vs. Unexpected Inflation in Sargent’s Conquest Model

Evidently, the Bank quickly discovers the folly of its ways, and inflation is quickly stabilized, *forever*. In the words of Sargent (1999), this would truly constitute a ‘triumph of the natural rate hypothesis’.

2.1.3. *Specification Searches.* Unfortunately, this propitious outcome is not the only possibility. Suppose a different researcher arrives at the Bank, who argues that the Bank’s model is misspecified because it omits supply shocks, which shift the Phillips Curve. In addition, he persuades the Bank to modernize its econometric methods. Rather than formal hypothesis testing, he argues that it is more reasonable, in light of misspecification concerns, to apply a model selection approach, based on Akaike’s Information Criterion. Also, he points out that it is crazy to repeatedly test the same hypothesis using a constant critical value. Instead, he recommends a ‘monitoring structural change’ approach, as developed by Chu, Stinchcombe, and White (1996).⁴ So let’s now suppose the Bank

⁴Presumably, the researcher has access to a time machine, as these methods were not developed until recently.

entertains *two* models, one being Sargent’s simple bivariate Phillips Curve, and the other being a supply-shock augmented model containing real oil prices as an additional regressor. The Bank selects between these models using an AIC statistic, but sticks with its status quo model as long as it remains within the ‘no structural change’ bounds determined by the monitoring procedures of Chu, Stinchcombe, and White (1996). In this case, rather than simulations, it is more interesting to see what would have happened in light of actual U.S. inflation history. Figure 4 contains the results⁵




FIGURE 4. Monitoring Structural Change in a Phillips Curve Model

The top panel displays the perceived slopes of the Phillips curve in the two models, i.e., the estimated inflation coefficients. The bottom panel then displays the resulting inflation target. Note that when the Bank’s model includes oil prices, the inflation target responds to the current price of oil. The response is positive since oil prices are estimated to have a negative effect on employment.

It turns out the Bank switches models four times: (1) In late 1966 it switches to the model with oil prices, even though oil prices were relatively flat until then. This likely reflects the combination of gradual negative trends in both unemployment and real oil prices, (2) At the end of 1969 it switches back to the static Phillips curve. This reflects a decline in the correlation between oil prices and unemployment, as oil remains flat, while unemployment rises gradually during the prior year, (3) In early 1975 the Bank switches again to the supply-shock augmented Phillips curve. Notice that the switch is not instantaneous, since the first shock occurs at the end of 1973. This reflects the adaptive nature of the Bank’s beliefs, combined with a relatively small gain parameter ($\eta = .015$). Interestingly, the Bank retains the supply-shock augmented Phillips curve model until mid-2000, when it goes back to the simple static Phillips curve. This reflects the fact that the (partial) correlation between oil prices and unemployment remains significant

⁵A few details: (1) Inflation and unemployment data, as well as parameter values, are the same as in Sargent (1999), (2) The Bank starts its experiment in June 1960, with a static Phillips Curve calibrated to data from the prior year, which yields initial conditions $\hat{\beta}_0 = 4.95$ and $\hat{\beta}_1 = -0.29$, (3) Oil prices are taken from the IMF database, line 76AAZ, and deflated using the CPI, line 64, (4) The monitoring thresholds are taken from Deshayes and Picard (1982), and are given by $\pm\sqrt{n}\log\alpha$, with $\log\alpha = -2$. This corresponds to an approximate (nominal) significance level of 13%. Due to the endogenous nature of the DGP, the actual size will be different. The results are not highly sensitive to α .

well into the 1980s, although of course during the 1980s things were moving in reverse, with oil prices and unemployment both falling.

What's interesting about Figure 4 is that it shows how the ability to engage in specification searches enables the Bank to sustain at least some belief in the Phillips curve. However, it is clear that by themselves oil prices cannot explain the rise and persistence of U.S. inflation during the 1970s. Instead, they mainly seem to explain why the Fed might have retained a low, but positive, inflation target throughout the 1980s and 90s. The failure to generate a rise in target inflation during the 1970s can be partly blamed on the low gain parameter. In Figure 4, the Bank does weight recent data more heavily, but not by much. One could argue, however, that the oil shocks were indeed just that, i.e., shocks, in which case a constant gain algorithm, designed to track slow parameter drift, is not really appropriate. Although there are certainly more sophisticated ways to allow the gain to vary (e.g., Kushner and Yin (1997)) Figure 5 shows what happens when the gain increases to $\eta = 0.5$ during the turbulent periods immediately following the two oil shocks, 1974-75 and 1980-81.

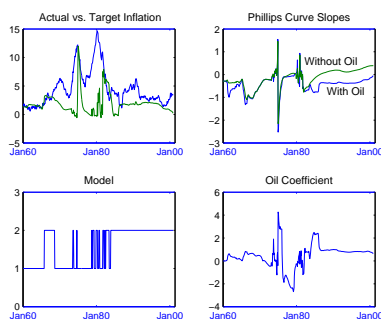


FIGURE 5. Model Selection, Oil Prices and a Variable Gain

Now we see a fairly significant jump in target inflation. Interestingly, it is the *combination* of a variable gain and model switching that generates the volatility in target inflation. In particular, notice that the slopes of the two Phillips curves are quite similar during the 1970s, and that the Bank only switches to the shock-augmented model for a few months during the first oil crisis, and switches back-and-forth quite rapidly during the second crisis. Still, when combined with the spike in the oil coefficient, this is enough to generate inflation movements that bear at least some resemblance to what actually occurred.

It is interesting to compare this account of recent U.S. inflation history to the account in Cogley and Sargent (2005), which attributes the Fed's delay in bringing down inflation to Bayesian model uncertainty. In their model, the Fed sticks to a high inflation policy because it risk-dominates the policy implied by the better-fitting model. In contrast, here the rise and persistence of inflation reflects the Fed's continued belief in the Phillips curve. This belief is sustained by the Fed's adaptive behavior, with respect to both parameters *and* models, which allows the Fed to rationalize what might otherwise be interpreted as a breakdown of the Phillips curve.

2.2. Model Validation in a Cobweb Model. Many macroeconomic models take the form of expectational difference equations. One prominent example is the famous cobweb model. This model has long been a useful laboratory for analyzing various issues in dynamic economics, first with constant coefficients adaptive expectations, then with rational expectations, then with adaptive least-squares learning, and most recently, with misspecified adaptive least-squares. We continue this tradition by using it to study model validation dynamics. In particular, we pursue an example studied by Evans and Honkapohja (2001, pgs. 318-20). They analyze the E-stability properties of so-called Restricted Perceptions Equilibria (RPE), in which agents (exogenously) omit relevant variables from their fitted models. In their analysis, *any* model can be a RPE, as long as its estimated coefficients adjust to account for the omitted variable bias. Here we allow agents to *test* their models, and ask whether some RPE are more resilient than others.

Consider then the following reduced form model,

$$p_n = \alpha E_{n-1} p_n + \beta_1 w_{1,n-1} + \beta_2 w_{2,n-1} + \varepsilon_n \quad (2.1)$$

where $(w_{1,n}, w_{2,n})$ are zero mean exogenous variables, and ε_n is an i.i.d. shock. Assume $w_{i,n}$ and ε_n are Gaussian. This model has a trivial Rational Expectations equilibrium,

$$p_n = \frac{1}{1 - \alpha} (\beta_1 w_{1,n-1} + \beta_2 w_{2,n-1}) + \varepsilon_n$$

From Evans and Honkapohja (2001), we also know that it has a pair of RPE. Let RPE_1 denote the RPE obtained by including $w_{1,n}$ and excluding $w_{2,n}$. Let RPE_2 denote the reverse case. We then have:

$$RPE_1 : \quad p_n = (1 - \alpha)^{-1} (\beta_1 + \Omega_{11}^{-1} \Omega_{12} \beta_2) w_{1,n-1} + \varepsilon_n \quad (2.2)$$

$$RPE_2 : \quad p_n = (1 - \alpha)^{-1} (\beta_2 + \Omega_{22}^{-1} \Omega_{12} \beta_1) w_{2,n-1} + \varepsilon_n \quad (2.3)$$

where Ω_{ij} are the elements of the second moment matrix, $\Omega = E(w w')$. Both RPE are E-stable if and only if $\alpha < 1$.

Now suppose an agent simultaneously entertains *both* models, continuously testing one against the other. For simplicity, suppose that he does not consider the correctly specified larger model.⁶ Further suppose the agent wants to guard against parameter drift within each model, so he adopts a constant gain recursive learning algorithm.

As it stands, this set-up is the same as in Branch and Evans (2006). Their goal is quite different, however. They posit a large collection of agents who randomly select between the two models, with weights determined by recent forecasting performance. They show that heterogeneity of expectations can persist even when agents pick the best model with probability converging to unity. In contrast, we posit a single agent who deliberately injects randomness into his model choice. This randomness does not reflect capriciousness or computational errors. Rather, it represents a strategic response to model uncertainty. In doing this, we borrow insights from the calibrated learning literature, which has shown that randomized forecasts are essential in delivering good forecasts in the presence of model uncertainty. (See, e.g., Foster and Vohra (1998)). Randomization prevents ‘nature’ from exploiting an agent’s ignorance of the data-generating process, and in this sense it

⁶Brock and Hommes (1997) consider a similar model, and show that interesting dynamics arise when agents can, with an additional cost, use the correct specification.

ensures a form of ‘robustness’. We shall show that randomization is also important in preventing agents from getting stuck in ‘bad’ Self-Confirming Equilibria.

Figure 6 shows what happens when an agent selects between the two models using a discrete-choice/logit function. Each model is recursively estimated using a constant gain learning algorithm. Recursive estimates of each model’s mean-squared forecast error are then used in the logit function to update the model weights. The parameters have been rigged to favor Model 1, in the sense that the variance of w_1 is greater than the variance of w_2 , so that Model 2 omits the more important variable.⁷




FIGURE 6. Model Selection in a Misspecified Cobweb Model

Not surprisingly, the agent uses Model 1 most of the time. Randomization leads to frequent, but short-lived, experiments with Model 2. However, occasionally (every 1000+ periods or so) a prolonged excursion to Model 2 takes place. These excursions represent the model’s escape dynamics. They can be analytically characterized using the tools of large deviations theory. The RPE of each model has a large deviations rate function attached to it. These rate functions summarize how difficult (ie., unlikely) it is to escape from each of the RPE. In the limit, as the update gain gets smaller, one model becomes dominate, and is used ‘almost always’. We use large deviations methods to identify robust Self-Confirming Equilibria. We show that robust SCE have relatively large rate functions.

3. A GENERAL FRAMEWORK

This section attempts to distill the lessons of the previous examples into a general analytical framework. This framework is sufficiently general to encapsulate most macroeconomic policy problems.

We study a decision maker who must solve a Linear Quadratic Regulator (LQR), but he doesn’t know the state transition equation. Traditional adaptive control methods are designed to handle coefficient uncertainty, but not model uncertainty. Recently developed

⁷Some more details: (1) w_1 and w_2 are assumed independent with $\sigma_{w_1}^2 = .45$ and $\sigma_{w_2}^2 = .35$, (2) $\sigma_\varepsilon^2 = .25$, (3) $\beta_1 = \beta_2 = 1.0$, (4) $\alpha = 0.5$, (5) The update gain is set to $\eta = .02$, and (6) The ‘choice intensity’ parameter of the logit function is set to 25.0.

robust control methods are designed to handle general forms of model uncertainty, *assuming* the decision maker starts with an exogenously specified reference model against which all uncertainty can be measured. (See Hansen and Sargent (2007b)). Bayesian methods could in principle be applied, but would require specification of high dimensional prior distributions. Moreover, Bayesian methods lead naturally to model averaging, and encounter well known difficulties when used to select models.⁸ Our approach is to assume the decision maker starts with a *model class*, and then tries to sort out, in real-time, among this class using traditional specification testing methods.⁹

3.1. Objective Function. Consider a decision maker who cares about an $s \times 1$ vector of state variables, x_t , and who believes he can influence them by choosing a $c \times 1$ vector of control variables, u_t . Preferences over these variables are ordered by the following Linear-Quadratic loss function,

$$E_n \sum_{j=0}^{\infty} \delta^j \{x'_{n+j} Q x_{n+j} + u'_{n+j} R u_{n+j}\} \quad (3.4)$$

where Q and R are $s \times s$ and $c \times c$ positive definite matrices, and δ is a discount factor. In what follows, it is important to keep in mind that the expectations operator here pertains to the beliefs of the decision maker. The probability measure associated with these beliefs may differ from the probability measure generating the data, and it may evolve over time.

3.2. Models. The decision maker interprets the mapping between u_n and x_n using a set of candidate models, \mathcal{M} , containing m elements. Each model is linear and Gaussian, and is described by a finite collection of unknown parameters. The models may be nested or non-nested. Our maintained assumption is that the decision maker believes there to be a single ‘best’ model within this class, and the goal is to find it, while at the same time balancing the ongoing pressures of meeting the control objective. An important limitation of our analysis is the assumption that the model class, \mathcal{M} , is fixed. The decision maker refines estimates of each model’s parameters, and revises his beliefs over time about the usefulness of each model, but by assumption, we do not allow the decision maker to reconsider his original specification of \mathcal{M} .¹⁰

We assume the decision maker entertains the possibility that the state variables of interest, x_n , are embedded within some larger system, which in addition to x_n , contains an $n \times 1$ vector of exogenous variables, z_n . It is possible these exogeneity restrictions are misspecified. In effect then, each model $i = 1, 2, \dots, m$ is defined by a set of s variable selection matrices, e_j^i $j = 1, 2, \dots, s$, specifying which of the elements of x_n and z_n it contains, and a set of s control selection matrices, $e_{u,j}^i$, specifying which of the control

⁸Brock, Durlauf, and West (2007) make progress in applying Bayesian methods to macroeconomic policy design.

⁹Interestingly, Gilboa, Postlewaite, and Schmeidler (2008) argue that classical, frequentist-based econometric methods are actually more in keeping with recent developments in decision theory than are Bayesian methods.

¹⁰Of course, this is only an issue if events seem to call for an expansion of the model class. By adjusting the weights, the decision maker can effectively discard (relatively) poorly performing models.

variables are relevant for each equation. If we then define the $(s + n) \times 1$ stacked vector $y_n = (x_n, z_n)$, we can express the elements of \mathcal{M} as follows

$$x_{jn}^i = \alpha_j^i e_j^i y_{n-1} + \gamma^i e_{u,j}^i u_{n-1} + \varepsilon_{jn}^i \quad i = 1, 2, \dots, m \quad j = 1, 2, \dots, s \quad (3.5)$$

As usual, there is no loss in generality in assuming a single lag of y_t in each state transition equation, as higher order lags can be accommodated by increasing the dimensionality of the state. What is essential, and in fact what is the defining feature of each model, is the assumption that the error terms, ε_{jn}^i , are i.i.d. and uncorrelated with the regressors. The specification test will be designed to check the accuracy of this assumption.

3.3. Data-Generating Process. We assume there is an underlying time invariant data-generating process (DGP). The agent is trying to learn about this process. What makes this problem difficult is that the DGP *responds* to the agent's model selection and control decisions. Of course, if the agent knew the model, he could follow the advice of Lucas (1976) and account for this endogeneity in the formulation of his policies. Without knowing the model, however, it is not at all clear how he should proceed. The danger is that he gets trapped in a 'bad' self-confirming equilibrium. One response would be to assume the agent is simply unaware of this danger. Our view is that policymakers are in fact often aware of feedbacks, but respond to them informally. Good policymakers know intuitively that is important to occasionally try new things. We model this process very simply by assuming that model selection is subject to some deliberate randomness.¹¹

Letting $\beta_n = (\alpha_{jn}^i, \gamma_{jn}^i)$ be the full vector of model coefficient estimates at time- n , we can capture the feedbacks from beliefs to the DGP by writing it as follows,

$$x_n = A_1(s_{n-1}, \beta_{n-1})x_{n-1} + B_1(s_{n-1}, \beta_{n-1})z_{n-1} + C_1(s_{n-1}, \beta_{n-1})u_n + D_1(s_{n-1}, \beta_{n-1})\varepsilon_{1n} \quad (3.6)$$

$$z_n = A_2(s_{n-1}, \beta_{n-1})z_{n-1} + B_2(s_{n-1}, \beta_{n-1})x_{n-1} + C_2(s_{n-1}, \beta_{n-1})u_n + D_2(s_{n-1}, \beta_{n-1})\varepsilon_{2n}$$

where s_n is an indicator for the model being used at time- n . The coefficient functions encoding the feedback from beliefs to the actual DGP are case specific, and can be quite complex and highly nonlinear. Fortunately, all we need is the following assumption,

Assumption 3.1. *For all β_n and s_n , the joint distribution of x_t and z_t is stationary and ergodic.*

Our analytical methods rely heavily on the ability to 'average out' fluctuations in x_t and z_t for given values of the model coefficients and the model indicator. They will not work without this assumption.

3.4. Model Updating. To the extent the decision maker is aware that his own model selection influences the DGP, he faces a difficult counterfactual - How would things have been different if instead a different model had been used in the past? Fitting a model to data that was generated while a *different* model was in use could produce very misleading inferences about the prospects of a given model. This counterfactual lies at the heart of macroeconomic policy formulation. Theory and deductive reasoning no doubt have a role to

¹¹Fudenberg and Levine (2008) discuss the importance of experimentation in macroeconomic policy. On the other hand, Cogley, Colacito, Hansen, and Sargent (2008) show that a concern for robustness in the face of model uncertainty may temper the motive to experiment.

play in resolving this counterfactual dilemma, but absent a formalization of this process, we assume our decision maker adopts the expedient of only updating models while they are being used. The previous track records of unused models still provide competitive pressure on the performance of the current model, but there is a sense in which these other models are ‘on ice’ while the current model is being used. What makes this work is our assumption of persistent experimentation. Eventually, the decision maker comes to have a fairly accurate picture of how each model performs.

Other than this occasional hibernation, models are updated in the standard way, using a (discounted) recursive least-squares algorithm. Unfortunately, writing this algorithm in the conventional way requires a little notation. Start by writing 3.5 as follows

$$\begin{aligned} x_{jn}^i &= (\alpha_j^i, \gamma_j^i) \begin{pmatrix} e_j^i & 0 \\ \mathbf{0} & e_{uj}^i \end{pmatrix} \begin{pmatrix} y_{n-1} \\ u_{n-1} \end{pmatrix} + \varepsilon_{jn}^i \\ &= \beta_j^i E_j^i \phi_{n-1} + \varepsilon_{jn}^i \end{aligned}$$

where E_j^i collects all the regressors in the j th equation of model i , and β_j^i collects all their coefficients. Then, if we place all s equations into the $1 \times s$ vector x_n^i , we can write

$$x_n^i = \beta^{i'} \Phi_{n-1}^i + \varepsilon_n^i \quad i = 1, 2, \dots, m$$

where $\Phi_{n-1}^i = \text{diag}[E_1^i \phi_{n-1}, E_2^i \phi_{n-1}, \dots, E_s^i \phi_{n-1}]$. Finally, if we define s_n^i as an indicator function for whether model i is being used at time- n , we have the update equations

$$\hat{\beta}_n^i = \hat{\beta}_{n-1}^i + \eta s_{n-1}^i (R_{n-1}^i)^{-1} \Phi_{n-1}^i (x_n^{i'} - \Phi_{n-1}^{i'} \hat{\beta}_{n-1}^i) \quad (3.7)$$

$$R_n^i = R_{n-1}^i + \eta s_{n-1}^i (\Phi_{n-1}^i \Phi_{n-1}^{i'} - R_{n-1}^i) \quad (3.8)$$

where η is a constant gain parameter, assumed to be common across equations and models. Note that no effort is made to exploit potential cross-equation error term correlations. The decision maker simply employs equation-by-equation least squares.

3.5. Self-Confirming Equilibria. It is important to notice that the x_n vector on the right-hand side of (3.7) corresponds to the actual law of motion given by (3.6), which depends on both the control and the estimation efforts of the decision maker. This makes the decision maker’s problem ‘self-referential’. It also makes the analysis of this problem difficult. As discussed in detail in Section 4, our approach is to exploit a time-scale separation between the evolution of the data, Φ_n , and the evolution of each model’s coefficient estimates, $\hat{\beta}_n^i$. The key concept when doing this is the notion of a ‘mean ODE’, which is obtained by taking the following four steps: (1) Substitute the actual law for x_n given by (3.6) into the parameter update equation in (3.7), (2) Freeze the coefficients at their current values, (3) Average over the stationary distribution of the ‘fast’ variables, Φ_n , which exists by Assumption 1, and (4) Form a continuous time interpolation of the resulting autonomous difference equation, and then obtain the mean ODE by taking limits as $\eta \rightarrow 0$. The only subtlety here arises from the model switching. In section 4 we show this switching occurs on a time scale that is even slower than the estimation time scale, so by itself this creates no special problems when obtaining the mean ODE. However, since we assume coefficients are only updated while a model is in use, the stationary distribution in the mean ODE is model specific, i.e., there is a separate distribution for each model.

For notational convenience, we write each model’s resulting mean ODE as follows:

$$\dot{\beta}^i = h^i(\beta^i) \tag{3.9}$$

Note, the vector function $h^i(\cdot)$ encodes both the least squares orthogonality conditions and the averaging over the (model specific) stationary distribution of the data.

We can now state the following important definition:

Definition 3.2. *A Self-Confirming Equilibrium (SCE) for model i consists of a fixed vector of regression coefficients, β^i , and a control feedback matrix, $F^i(\beta^i)$, which satisfy the following two conditions:*

- (a): *The coefficients are a stationary point of the mean ODE, $h^i(\beta^i) = 0$.*
- (b): *The control policy $u_n^i = -F^i(\beta^i)y_{n-1}$ solves the LQR in (3.4) for model i*

At this level of generality there is not much to be said about existence of uniqueness of SCE. Since our analysis focuses on selection among models and their associated SCE, it doesn’t make much sense without the following assumption:

Assumption 3.3. *Each model $i = 1, 2, \dots, m$ has a unique, asymptotically stable Self-Confirming Equilibrium, β^{*i} .*

Uniqueness is not essential. It just simplifies the analysis. Stability is not essential either, but without it the problem becomes uninteresting. In a sense, we are conducting a kind of ‘meta-stability’ analysis, and asking which, among a set of SCE, will be the most resilient to specification testing. If a model’s SCE is unstable, then it will not even survive within its own learning dynamics.

3.6. Model Validation. We have been purposely vague about the details of specification testing. That’s because there is no single best way to validate a model. The right approach depends on what the model is being used for, and the nature of the relevant alternatives. Since we begin with an explicit model class, \mathcal{M} , a natural approach would be to run a sequential ‘horse race’ among the models, and use the model that appears to offer at each point in time the best approximation to the unknown DGP, as measured by an AIC or BIC statistic. However, this can in general produce frequent, destabilizing, model switching. In fact, this is one of the main arguments in favor of model averaging. Although there are undoubtedly good statistical reasons to average across models, especially in a forecasting context, we doubt its effectiveness in a control context. From a control standpoint, it is advantageous to have a simple model. Although it would of course be desirable to make these model switching and control costs explicit, we simply assume the decision maker needs to select a single, relatively parsimonious, model. He does this by engaging in the following two step validation process - First, the currently used model is viewed as a reference point, or ‘null hypothesis’, and will continue to be used unless sufficient evidence mounts against it. This anchoring on the current model reflects implicit model switching costs. Second, if it turns out the current model is rejected, then a new model is (randomly) selected from the model class, \mathcal{M} . As noted earlier, the random element in the selection process is motivated by strategic considerations. It’s a form of experimentation in response to unstructured model uncertainty. However, it can also be interpreted as a form of model averaging, but with the important difference that this averaging, or hedging, only takes

place *ex ante*. After the fact, since the model must be used to solve a control problem, the decision maker commits himself to whichever model was drawn.

Although our approach is sufficiently general that it could accommodate a variety of alternative specification tests, we opt for a Lagrange Multiplier (LM) approach. LM tests can be interpreted as Likelihood Ratio tests against localized alternatives (Engle (1984)), or as first-order approximations of KLIC statistics. Their defining feature is that they are based solely on estimation of the null model, and do not require specification of an explicit alternative. For this reason, they are often referred to as *misspecification* tests, rather than specification tests.

In addition to being widely applied within the conventional econometrics literature, LM testing has also received some attention within the literature on Stochastic Recursive Algorithms (SRA). Chapter 5 in Benveniste, Metivier, and Priouret (1990), in particular, outlines a recursive validation procedure based on LM testing principles. Their method is based on the observation that the innovation in a typical SRA is proportional to the score vector.¹² Essentially then, what is being tested is the significance of the algorithm’s update term.

Our approach is similar to that of Benveniste, Metivier, and Priouret (1990), except we take the additional step of asking what happens in the event of model rejection. A common criticism of misspecification testing is that it fails to address the question of how to respond to test rejections. Our agent responds by going back to his original model class and looking for a better model.

Figure 7 shows what would happen if the Central Bank were to apply a recursive LM testing strategy in the context of Sargent’s (1999) *Conquest* model.¹³




FIGURE 7. Sequential LM tests in Sargent’s Conquest Model

It seems clear that during escapes the Central Bank would have cause to doubt the specification of its model. One should bear in mind, however, that this is suggestive

¹²Breusch and Pagan (1980) noted this as well, and argued that it could be exploited to simplify the computation of conventional (nonrecursive) LM tests.

¹³Some details: (1) Let $x_n = (1, \pi_n)$ be the regression vector, R_n be its second moment matrix, ξ_n be the time- n model residual, and let $\hat{\sigma}_n^2 = \hat{\sigma}_{n-1}^2 + \eta(\xi_n^2 - \hat{\sigma}_{n-1}^2)$, (2) The bottom panel of Figure 7 then reports the recursively estimated statistic, $\mathcal{L}_n = \mathcal{L}_{n-1} + \eta[(x_n' \xi_n) R_n^{-1} (x_n \xi_n) / \hat{\sigma}_n^2 - \mathcal{L}_{n-1}]$.

evidence, at best. First, although one can obtain distributional results in traditional settings, the application here departs from these settings in some important respects (e.g., endogenous data and constant gain updating). Second, as noted by Hansen (1992), the implicit alternative in this kind of test can be interpreted as being one of random walk parameter drift. However, as Sargent and Williams (2005) note, random walk parameter drift is the implicit *null* hypothesis behind constant gain learning algorithms. Applying LM tests to a model that is updated with a constant gain learning algorithm should therefore be designed to detect coefficient changes that are more rapid than what is implied by the algorithm's gain parameter. This can be accomplished most easily by subtracting a small positive constant from the squared residuals when computing the test statistic. In terms of the underlying model, what is then being tested is stability of the covariance matrix in the coefficient drift process.

Before proceeding, it will be useful to make explicit the connection between our model update equations and a traditional scoring algorithm by rewriting equation (3.7) as follows:

$$\hat{\beta}_n^i = \hat{\beta}_{n-1}^i + \eta s_{n-1}^i (R_{n-1}^i)^{-1} \Lambda_n^i \quad (3.10)$$

where Λ_n^i is defined as model i 's estimated vector of (unscaled) Gaussian scores

$$\Lambda_n^i = \Phi_{n-1}^i (x_n - \Phi_{n-1}^{i'} \hat{\beta}_{n-1}^i)$$

Model rejections will be triggered by excessively large fluctuations in Λ_n^i . Basing rejections on the vector Λ_n^i involves the following assumption:

Assumption 3.4. *Models are evaluated and tested in their entirety. The decision maker does not attempt to 'isolate' the source of any rejection.*

Of course, in terms of its descriptive accuracy, this is a highly debatable assumption. In practice, forecasters often do try to diagnose the source of model breakdowns. That's because in practice forecasters do not begin with a preconceived model class. Instead, the model class itself evolves in response to this diagnostic process. Unfortunately, modeling this open-ended process is much more difficult, so we stick to Assumption 3.4.

3.7. Model Selection. As noted earlier, a common criticism of misspecification tests is that they fail to specify an alternative. That's only partly true here. Our model validation process consists of two steps, where the first step consists of a misspecification test, and the second step specifies what happens in the event of a rejection. Our reasons for not combining these two steps into a single unified process of testing against explicit alternatives was previously discussed in section 3.6.

The model selection step of the validation process has two components: (1) The decision maker must specify a metric for model comparison. Once the current reference model is rejected, models must be compared. These metrics can then be normalized and used to construct a set of *model weights*. (2) Given the model weights, the decision maker must then construct a probability distribution across models. Earlier we motivated random selection on the basis of the agent's own strategic interests. It can also be motivated purely on the basis of technical considerations. Since \mathcal{M} is a finite set, random model selection can be interpreted as a 'convexification' of the constraint set.

Model weights are assumed to be based on model fit. There are, of course, well known dangers to basing model selection on model fit. The Phillips Curve example in Section 2

provides one example. Kocherlakota (2007) provides other examples. Since the models are being used to solve a control problem, a natural metric would be to base the weights on each model's implied value function. Given the LQR structure this would be computationally straightforward, and could be formulated as a 'Q-Learning' problem (Kushner and Yin (1997)). However, we believe there is some descriptive appeal to basing model selection on model fit. In practice, the tasks of model construction and control are often decentralized. A research department constructs models and gives advice to policymakers. Policymakers then decide on a control action. Although regular consultations take place, it is often the case that the goals and constraints of policymakers are imperfectly known to the research department. In this case, basing model selection on model fit makes more sense.

In general, when model comparisons are based on fit, it is essential to somehow account for model complexity. Otherwise there would be a natural tendency to use the 'biggest' model. In traditional one-sample, retrospective, settings this is typically accomplished by adding a 'penalty' term to the likelihood. There are a plethora of model selection criteria, distinguished by the details of this penalty term (see, e.g., Burnham and Anderson (2002) for a textbook discussion). Fortunately, it turns out that the real-time nature of the problem here works to our advantage. Since models are being used to forecast out of sample, they naturally embody an overfitting penalty. Overly parameterized models will produce high variance forecasts. Underparameterized models will produce low variance, but biased, forecasts. Basing model selection on ex post realized mean-squared error provides a natural resolution of this bias/variance trade-off.¹⁴

In keeping with the constant gain coefficient updating in equation (3.7), we assume constant gain updating of covariances and mean-squared errors. If we let $\nu_n^i(\beta_{n-1}^i)$ be the $s \times 1$ vector of time- n residuals from model i , and let ω_n^i be a measure of model i 's (lack of) fit, we have the following recursion for the weights:

$$\omega_n^i = \omega_{n-1}^i + \eta s_{n-1}^i (\nu_n^i(\beta_{n-1}^i)' \nu_n^i(\beta_{n-1}^i) - \omega_{n-1}^i) \quad (3.11)$$

Notice that, as with coefficient estimates, model weights, ω_n^i , are only updated while a model is in use (i.e., $s_n^i = 1$).

The second step is to now specify a mapping, $\Pi(\omega) : \mathbb{R}_+^s \rightarrow \Delta^{s-1}$, taking the model weights into a probability distribution over the models. Although any continuous function that has full support and satisfies the obvious monotonicity conditions will do, it is useful to be specific. The following logit function will be assumed:

$$\pi_n^i = \frac{e^{-\rho \omega_n^i}}{\sum_{j=1}^m e^{-\rho \omega_n^j}} \quad (3.12)$$

where $\rho > 0$ is a free parameter determining how much noise there is in the model selection process.

¹⁴Basing model selection on one-step ahead forecasts from sequentially estimated models is sometimes called the 'prequential plug-in' approach, and has been forcefully advocated by Dawid (1984). As noted by Hansen and Yu (2001), prequential plug-in embodies an implicit $(k/2) \log(n)$ penalty term, of the sort that arises from Bayesian and Minimum Description Length approaches to (retrospective) model selection.

4. ANALYSIS

Our goal is to understand the dynamics of macroeconomic model selection. Long ago, in response to the Lucas Critique, Sims (1982) argued that macroeconomic policymakers do not need to know the correct model of the economy in order to formulate good policy. According to Sims, as long as policy is flexible and adaptive, policymakers will learn to do the right thing.¹⁵ Due to the endogeneity of the DGP, it is not at all obvious whether this is indeed true. Sargent (1999) and Sargent and Williams (2005) study this question in settings where models are fixed, but parameters can vary. Their findings point to the importance of priors concerning parameter drift. Our analysis pursues the same questions, but in a setting with multiple models. Here there is greater scope for getting stuck in ‘bad’ Self-Confirming Equilibria. As in Fudenberg and Levine (2008), we are interested in studying the role of experimentation in selecting favorable SCE.

4.1. Representation as a Stochastic Recursive Algorithm (SRA). We have purposely tried to stay as close as possible to the standard SRA framework.¹⁶ These models feature an interplay between beliefs and outcomes. Our model validation framework features these same elements, but incorporates model testing and selection dynamics as well. It is useful to begin by collecting together the model’s equations:

We first have a set of model update equations,

$$\hat{\beta}_n^i = \hat{\beta}_{n-1}^i + \eta s_{n-1}^i (R_{n-1}^i)^{-1} \Lambda_n^i \quad (4.13)$$

$$\Lambda_n^i = \Phi_{n-1}^i (x_n - \Phi_{n-1}^{i'} \hat{\beta}_{n-1}^i) \quad (4.14)$$

$$R_n^i = R_{n-1}^i + \eta s_{n-1}^i (\Phi_{n-1}^i \Phi_{n-1}^{i'} - R_{n-1}^i) \quad (4.15)$$

Through feedback, these help to determine the actual DGP,

$$x_n = A_1(s_{n-1}, \beta_{n-1})x_{n-1} + B_1(s_{n-1}, \beta_{n-1})z_{n-1} + C_1(s_{n-1}, \beta_{n-1})u_n + D_1(s_{n-1}, \beta_{n-1})\epsilon_{1n} \quad (4.16)$$

$$z_n = A_2(s_{n-1}, \beta_{n-1})z_{n-1} + B_2(s_{n-1}, \beta_{n-1})x_{n-1} + C_2(s_{n-1}, \beta_{n-1})u_n + D_2(s_{n-1}, \beta_{n-1})\epsilon_{2n}$$

Model weights evolve according to the recursion

$$\omega_n^i = \omega_{n-1}^i + \eta s_{n-1}^i (\nu_n^i (\beta_{n-1}^i)' \nu_n^i (\beta_{n-1}^i) - \omega_{n-1}^i) \quad (4.17)$$

and in the event the current model is rejected, models are selected according to the probability distribution

$$\pi_n^i = \frac{e^{-\rho \omega_n^i}}{\sum_{j=1}^m e^{-\rho \omega_n^j}} \quad (4.18)$$

Models are tested by forming the recursive LM test statistic

$$\theta_n^i = \theta_{n-1}^i + \eta s_{n-1}^i [\Lambda_n^i \Sigma_{in}^{-1/2} (R_n^i)^{-1} \Sigma_{in}^{-1/2} \Lambda_n^i - \theta_{n-1}^i] \quad (4.19)$$

where Σ_{in} is an $s \times s$ diagonal matrix containing the squared residuals from model i . The model indicators, s_n^i , evolve as an m -state Markov Chain. Let $p_n \in \Delta^{m-1}$ be the time- n

¹⁵See, in particular, pg. 117 of Sims (1982).

¹⁶Benveniste, Metivier, and Priouret (1990) and Evans and Honkapohja (2001) contain good textbook treatments of SRA methods.

probability distribution over models, and let \mathcal{P}_n be an $m \times m$ state transition matrix, where $\mathcal{P}_{ij,n}$ is the time- n probability of switching from model i to model j . Model selection dynamics can then be represented as follows

$$p'_{n+1} = p'_n \mathcal{P}_n \quad (4.20)$$

The diagonal elements of \mathcal{P}_n are given by

$$\text{Prob}[\theta_n^i \leq \bar{\theta}(n)] + \text{Prob}[\theta_n^i > \bar{\theta}(n)] \cdot \pi_n^i \quad (4.21)$$

and the off-diagonals are given by

$$\text{Prob}[\theta_n^i > \bar{\theta}(n)] \cdot \pi_n^j \quad (4.22)$$

where $\bar{\theta}(n)$ are boundary functions calibrated to achieve a given (nominal) test size.

4.2. Time-Scale Separation. Equations (4.13) - (4.20) constitute a high-dimensional system of nonlinear stochastic difference equations. Shedding any kind of analytical light on this system would seem to be a hopeless task. The key to making the system tractable is the application of so-called ‘singular perturbation’ methods, which exploit the fact that subsets of the variables evolve on different time-scales. By appropriately averaging over subsets of the variables, we can simplify the analysis to one of studying the (weak) interactions between smaller subsystems, each of which can be studied in isolation.

We shall show that model validation dynamics feature a hierarchy of *three* time scales. The state and control variables evolve on a ‘fast’, calendar time-scale. The coefficients of each model evolve on a ‘slow’, model revision time-scale, where each unit of time corresponds to $1/\eta$ units of calendar time. Finally, model switching occurs on a ‘really slow’, large deviations time-scale, where each unit of model time corresponds to $\exp[S^*/\eta]$ units of coefficient time, where S^* is a model specific ‘rate function’, summarizing how difficult it is to escape from each model’s self-confirming equilibrium. This hierarchy of time-scales greatly simplifies the analysis of model validation, as it permits us to focus separately on the problems of control, model revision, and model selection. The novel aspect of our analysis is the ultimate, large deviations time scale. It involves rare but recurrent Markov switching among the finite set of models, each with coefficients fixed at their self-confirming values, and with the underlying data fixed at the mean of a model specific invariant distribution. In other words, we are going to replace the above time-varying Markov transition matrix, \mathcal{P}_n , with a *constant*, state-independent, transition matrix, $\bar{\mathcal{P}}$, with elements determined by the large deviations properties of each of the models. In the spirit of Kandori, Mailath, and Rob (1993), it will turn out that as $\eta \rightarrow 0$ the stationary distribution across models will collapse onto a single model.

4.3. Mean ODE Approximation of Model Revisions. We need to characterize the dynamic interactions among three classes of variables: (1) The state and control variables that appear as regressors within each model, Φ_n , (2) The coefficient estimates, β_n , and (3) The set of model indicators, s_n . We start in the middle, with the coefficient estimates. Their dynamics can be approximated by *averaging* over the Φ_n variables, while holding *fixed* the model indicators. Assumption 3.1 assures us that the averaging is well defined. Since updating occurs only while models are in use, we have a *system* of model specific ODEs. This implies, for example, that when obtaining the ODE approximation of model

i 's coefficients, the Φ_n data are averaged out with respect to the invariant distribution of 3.6 obtained by setting $s_{kn} = 1$ for $k = i$ and $s_{kn} = 0$ for $k \neq i$.

Model switching does create one subtlety not present in traditional applications of ODE approximations. As noted earlier, when models are rejected and the decision maker switches to a different model, the rejected model goes into a hibernation state, and will not be updated again until it is recalled (which occurs with probability one, given our experimentation assumption). It is essential that when a model is reawakened, its coefficients lie within the basin of attraction of its SCE. Thus we have:

Assumption 4.1. *Model rejections occur within the basin of attraction of each model's Self-Confirming Equilibrium*

Model rejections occur in response to the escape dynamics of each model. Escapes must overcome the model's mean dynamics, which attempt to suck the coefficients back to the SCE. They are driven by very special sequences of shocks. When a model is rejected in the midst of an escape, this shock sequence is short-circuited. Later, when the model is recalled to duty, new shocks and new initial conditions will have been realized, and all that will remain from its earlier state (with very high probability) is the mean dynamics. Hence, the coefficients must again wait for another opportunity to escape.

What the mean ODE is approximating then is the right tail portion of the return back to the SCE from the point of rejection (which, with very high probability, occurs at the same point within the model's parameter space).

Proposition 4.2. *Let \tilde{n} represent reinitialized discrete time in a model's most recent epoch of use (so that $\tilde{n} = 0$ corresponds to the date when a model begins to be used). Let $\beta_\eta^i(t) = \beta_{\tilde{n}}^i \quad \forall t \in [\eta\tilde{n}, \eta(\tilde{n} + 1))$ be a continuous-time interpolation of this reinitialized coefficient sequence. Then, as $\eta \rightarrow 0$, $\beta_\eta^i(t)$ converges weakly to the solution of model i 's mean ODE*

$$\dot{\beta}^i = h^i(\beta^i)$$

conditioned on the event $s_{kn} = 1$ for $k = i$ and $s_{kn} = 0$ for $k \neq i$.

Proof. See Appendix A. □

This result can be interpreted as a function space analog of the Law of Large Numbers. It implies that on any finite time interval the path of the interpolated process $\beta_\eta^i(t)$ closely shadows the solution of the ODE with arbitrarily high probability as $\eta \rightarrow 0$. We can also obtain an analog of the Central Limit Theorem by studying the fluctuations of $\beta_\eta^i(t)$ around this path. To do this, write the update equations for model i using the shorthand notation

$$\beta_n^i = \beta_{n-1}^i + \eta s_{n-1}^i H(\beta_{n-1}^i, \Phi_{n-1}^i, x_n)$$

so that $h^i(\beta^i) = E\{H(\beta^i, \Phi^i(\beta^i), x(\beta^i))\}$, where expectations are taken with respect to the stationary distribution of (3.7) with β_n^i fixed at β^i . Next, if we define the scaled difference between the interpolated process and the mean ODE

$$U_\eta^i(t) = \frac{\beta_\eta^i(t) - \beta^i(t)}{\sqrt{\eta}}$$

we can state the following result

Proposition 4.3. *Conditional on the event that model i continues to be used, as $\eta \rightarrow 0$ $U_\eta^i(t)$ converges weakly to the solution of the stochastic differential equation*

$$dU(t) = h_\beta^i(\beta^i(t))U(t)dt + \mathcal{R}^{1/2}(\beta^i(t))dW$$

where $h_\beta^i(\cdot)$ is the Jacobian of $h^i(\cdot)$ and $\mathcal{R}(\cdot)$ is the stationary long-run covariance matrix with elements

$$\mathcal{R}_{ij}(\beta) = \sum_{k=-\infty}^{\infty} \text{cov}[H_i(\beta, \Phi_k(\beta), x_k(\beta)), H_j(\beta, \Phi_0(\beta), x_0(\beta))]$$

Proof. In addition to the fundamental Assumption 3.1, the proof relies on several regularity conditions, e.g., linear growth and Lipschitz continuity restrictions on $H(\cdot)$, and assumptions concerning the derivatives of $h^i(\cdot)$. Given our purposes, we simply make the required assumptions. See, e.g., Evans and Honkapohja (2001) for details. \square

4.4. Large Deviations Approximation of Model Switching. Propositions 4.2 and 4.3 describe the *average* behavior of each model's coefficient estimates. Both are conditioned on a model's continued use. Eventually, however, for a fixed $\eta > 0$, the coefficient estimates will wander a significant distance from the SCE. We have in mind a situation where this potentially triggers a model switch. These switches are *rare*, in the sense that they occur in response to tail events in the model revision process. We must now characterize these tail events. We do this using the tools of large deviations (LD) theory.

The analysis consists of four main steps. First, using results from Cho, Williams, and Sargent (2002) and Williams (2001), we establish that each model's sequence of coefficient estimates satisfy a Large Deviations Principle. Second, we use the Contraction Principle to link the LD properties of the coefficient estimates to the LD properties of the LM test statistics. Third, we use the LD properties of the test statistics to construct a homogeneous Markov Chain approximation of the model selection process. Finally, using this approximation, we characterize the limiting model distribution, and identify a 'dominant' model in terms of its LD rate function.

We begin with a definition

Definition 4.4. *Let E be a separable Banach space. Suppose $\mathbf{S}_n, n > 0$ are E -valued random variables. It is said that $\{n^{-1}\mathbf{S}_n\}$ satisfies a Large Deviations Principle if there is a lower semicontinuous rate function $I : E \rightarrow [0, \infty]$, with compact level sets $I^{-1}([0, a])$ for all $a > 0$, such that*

$$\liminf_{n \rightarrow \infty} n^{-1} \log P(n^{-1}\mathbf{S}_n \in A) \geq - \inf_{x \in A} I(x)$$

for all open subsets $A \subset E$, and

$$\limsup_{n \rightarrow \infty} n^{-1} \log P(n^{-1}\mathbf{S}_n \in B) \leq - \inf_{x \in B} I(x)$$

for all closed subsets $B \subset E$

In our setting, \mathbf{S}_n will either be defined as a sequence of coefficient estimates, or as a sequence of test statistics, with E then corresponding to the relevant path space. The crucial object here is the rate function, $I(x)$. Definition 4.4 shows precisely the sense in which large deviation events are rare, i.e., their probability declines *exponentially* with n ,

and the rate function plays the role of a scale factor in this decline. If one process has a uniformly larger rate function than another, the relative frequency of its escapes will vanish. This will be of some importance when characterizing the limiting properties of model validation dynamics.

Unfortunately, in practice there is often a big difference between proving the existence of a rate function and actually finding it. Fortunately, in Linear-Quadratic Gaussian settings like ours, we can rely on the results of Williams (2001) and CWS (2002), which provide a complete characterization of the LD rate function in these settings. We simply adapt their results to our notation.

Proposition 4.5. *Given Assumption 3.1, and assuming tightness of the sequence of coefficient estimates (enforced if necessary by a projection facility), for each model i there exists a matrix-valued function $Q^i(\beta^i, R^i)$ such that β_n^i satisfies a Large Deviations Principle with a rate function, \bar{S}^i , given by the solution of the following variational problem:*

$$\bar{S}^i = \inf_{\dot{v}} \frac{1}{2} \int_0^t \dot{v}(s)' Q^i(\beta^i(s), R^i(s))^{-1} \dot{v}(s) ds$$

subject to

$$\begin{aligned} \dot{\beta}^i &= (R^i)^{-1} \bar{g}^i(\beta^i) + \dot{v} \\ \dot{R}^i &= \bar{M}^i(\beta^i) - R^i \\ \beta^i(0) &= \beta^{*i}, R^i(0) = \bar{R}^i, \beta^i(t) \notin G \quad \text{for some } 0 < t < T \end{aligned}$$

where $\bar{g}^i(\beta^i)$ is defined by $h^i(\beta^i) \equiv (R^i)^{-1} \bar{g}^i(\beta^i)$ and $\bar{M}^i(\beta^i)$ is the second moment matrix of Φ_n^i when β_n^i is fixed at β^i .

Proof. The required regularity conditions are similar to those in Proposition 4.3. See Williams (2001) for details. \square

The striking thing about Propositions 4.2 and 4.5 is that together they characterize the sample paths of a nonlinear stochastic dynamic process in terms of the solutions of two *deterministic* differential equations; one characterizing the mean dynamics and the other characterizing the escape dynamics. The key object here is the weighting matrix, $Q^i(\cdot)$. It is determined by the covariance matrix of the Least Squares orthogonality conditions, and therefore depends on *fourth* moments. Williams (2001) shows that it can be computed by solving a system of matrix Lyapunov equations. Since the Least Squares orthogonality conditions are essentially the model's score vector, $Q(\cdot)$ will be closely related to the variance of the specification test statistic.

Solution of the large deviations control problem in Proposition 4.5 involves a minimization over points on the boundary, ∂G , of the parameter space. It turns out that with overwhelming probability the large deviations sample path hits the boundary at a unique point. This fact could be used to calculate efficient test statistics based directly on fluctuations in the coefficient estimates. However, an even better approach is to base inferences on the sequence of estimated scores. Under the null, these will behave as innovations, and therefore will more clearly reveal alternatives featuring breaks or other structural

changes.¹⁷ Hence, we need to translate the LD results for the coefficients into LD results for the LM test statistics in equation (4.19).

To do this, define the function, $\mathcal{Z}^i : \mathbb{R}^{k_i} \rightarrow \mathbb{R}_+$, where k_i is the number of variables in model i , as the score function, $\mathcal{Z}^i(\beta_n^i) = \Lambda_n^{i'} \Sigma_{in}^{-1/2} (R_n^i)^{-1} \Sigma_{in}^{-1/2} \Lambda_n^i$. Next, define the continuous-time interpolation of the LM test statistic

$$\theta_\eta(t) = \theta_\eta(0) + \eta \sum_{i=0}^{\lfloor t/\eta \rfloor} [\mathcal{Z}(\beta(i)) - \theta_\eta(i)]$$

and then define $\theta(t) = \lim_{\eta \rightarrow 0} \theta_\eta(t)$ as its limit. Now, observe that any given continuous-time path, $\beta(t)$ for a model's coefficients induces a corresponding path for its test statistic. (Remember that the state variables are being averaged out here). Define this mapping, $F : C[0, T] \rightarrow C[0, T]$, as follows

$$\begin{aligned} \theta(t) &= \theta(0) + \int_0^t [\mathcal{Z}(\beta(s)) - \theta(s)] ds \\ &\equiv F(\beta) \end{aligned} \tag{4.23}$$

We now have the following result

Proposition 4.6. *Each model's LM test statistic process, $\theta(t)$, has a locally stable equilibrium at $\theta^* = \mathcal{Z}(\beta^*) = 0$, and it satisfies a large deviations principle with rate function given by*

$$I(\theta) = \inf_{\{\beta: \theta = F(\beta)\}} \frac{1}{2} \int_0^t [\dot{\beta} - R^{-1} \bar{g}(\beta)]' Q(\beta(s), R(s))^{-1} [\dot{\beta} - R^{-1} \bar{g}(\beta)]$$

subject to $\theta(t^e) \notin G^\theta$ for some $0 < t^e < T$, where ∂G^θ defines a rejection threshold.

Proof. The stability of θ^* is clear from inspection of (4.23). The proof of the LDP is based on the following useful result from the theory of large deviations (see Dembo and Zeitouni (1998), p. 126) \square

Theorem 4.7. (Contraction Principle) *Let X and Y be Hausdorff topological spaces and $f : X \rightarrow Y$ a continuous function. Consider a (good) rate function $S : X \rightarrow [0, \infty]$.*

(a): *For each $y \in Y$ define*

$$S'(y) = \inf\{S(x) : x \in X, y = f(x)\}$$

Then S' is a (good) rate function on Y , where the infimum over the empty set is taken as ∞

(b): *If S controls the LDP associated with a family of probability measures μ_ε on X , then S' controls the LDP associated with the family of probability measures $\mu_\varepsilon \circ f^{-1}$ on Y .*

Loosely speaking, the contraction principle tells us that large deviations principles are preserved by continuous mappings. Of course, depending on the properties of f , the rate function S' might be quite different from the rate function S , so the large deviation

¹⁷Benveniste, Metivier, and Priouret (1990) emphasize this point. See p. 182.

properties of x and y themselves (e.g., escape times and escape routes) might be quite different. However, the contraction principle provides a means for translating between the two. To apply this theorem, we must establish that $F(\beta)$ is continuous.

Lemma 4.8. *If $\mathcal{Z}(\beta_n)$ is Lipschitz continuous, then $F(\beta(t))$ is continuous.*

Proof. See Appendix B □

Lipschitz continuity of $\mathcal{Z}(\beta_n)$ is a rather weak assumption. It will be satisfied as long as second moments remain bounded and there are no linear dependencies among the regressors.

According to Proposition 4.6, corresponding to any escape route for a model's coefficients there is an escape route for its LM test statistic. The idea is that as $\beta(t)$ departs from its SCE, $\theta(t)$ approaches an appropriately defined rejection threshold. In traditional settings, where the data are exogenous, we could safely assume that as the $\beta(t)$ wanders away from β^* (e.g., in Euclidean distance), then $\theta(t)$ monotonically increases. However, when the DGP is endogenous, it's conceivable that the variance of Λ falls by enough to make $\mathcal{Z}(\beta)$ decline with β . To get a sensible test boundary, we must rule this out

Assumption 4.9. *For each model and for all $t > 0$, $\mathcal{Z}^i(\beta^i(t))$ is a strictly increasing function of $\beta^i(t)$.*

The following figure provides a visualization of the sort of construction we have in mind




FIGURE 8. The Contraction Principle

In the applications of CWS (2002) and Williams (2001), ∂G is a free ‘parameter’. Increasing the radius of ∂G simply scales up \bar{S}^* . In our model validation framework, ∂G must be calibrated to achieve a desired (nominal) significance level. Remember the implicit null here embodies *slow* parameter drift. We only want to reject in response to what appears to be excessively ‘large’ and excessively ‘fast’ changes in coefficient estimates. Hence, we must actually use the Contraction Principle in both directions.

The analysis so far has exploited a time-scale separation between the data and each model's coefficient estimates. We've studied the evolution of a model's coefficients by

averaging out the dynamics of the state variables. Everything has been conditional on s_{it} , i.e., the current model. The next step in our analysis exploits a different kind of time-scale separation; namely, between the coefficient estimates and the frequency of model switching. Since model rejections are rare in the large deviations sense, we can now average out the dynamics in $\beta^i(t)$ and focus on switches *between* models. To do this we must define a new logarithmic time-scale, $\tau = \eta \log(t)$, where τ can be interpreted as the time-scale over which model switching occurs. In other words, each unit of model switching time, τ , corresponds to $\exp[\eta^{-1}]$ units of model revision time. Large deviation events only become ‘visible’ on this scale. Over this length of time we can average over the excursions that $\beta(t)$ takes away from the SCE, and fix its value at β^* (its long-run average), just as we fixed the values of the state variables at their stationary equilibrium values when studying the dynamics of $\beta(t)$. In fact, to obtain the necessary averaging for *all* models, we must actually employ the time-scale $(\eta/\bar{S}_{\max}) \log(t)$, where \bar{S}_{\max} is the largest LD rate function among all the models. The key to making this work is the following,

Proposition 4.10. *The mean escape time to ∂G is the same for all $\beta \in \overset{\circ}{G}$. Likewise, the mean escape time to ∂G^θ is the same for all $\theta \in \overset{\circ}{G}^\theta$.*

Proof. The proof follows from the uniqueness of the escape routes. \square

Intuitively, any β other than the exit point will, with overwhelming probability, converge back to the SCE before it makes another escape attempt. This generates a sort of state independence that enables model switching to be described as a homogeneous Markov Chain. As when studying the evolution of a model’s coefficients, we start by defining a continuous-time interpolation of the discrete distribution over the models, p_n . Over short horizons, the transition probability matrix, \mathcal{P}_n , of this Markov Chain is quite complex (see eqs. (4.21)-(4.22)). Our goal is to simplify this matrix by applying singular perturbation methods. Define the continuous-time interpolation of p_n as usual, i.e., $p^\eta(t) = p_n$ for $t \in [\eta n, \eta(n+1))$. Next, use the change of variables $\tau = (\eta/\bar{S}_{\max}) \log(t)$, and consider the rescaled process, $p^\eta(\tau)$. This process can be characterized as an m -state homogeneous Markov Chain.

Proposition 4.11. *Assume $\forall i \in \{1, 2, \dots, m\}$ that $\theta^i(t)$ is calibrated to reject during escapes of Model i . Assume $\pi^i(t) \in [\underline{a}, \bar{a}] \forall i, t$, where $\underline{a} > 0$ and $\bar{a} < 1$. Then as $\eta \rightarrow 0$, $p^\eta(\tau)$ converges weakly to a homogenous m -state Markov Chain with generator Q ,*

$$q_{ij} = \pi_j^* e^{(\bar{S}_{\max} - \bar{S}^i)/\eta} \quad q_{ii} = - \left(\sum_{j \neq i}^m \pi_j^* \right) e^{(\bar{S}_{\max} - \bar{S}^i)/\eta}$$

which possesses a unique invariant distribution

$$\bar{p}_i^\eta = \frac{\pi_i^* e^{\bar{S}^i/\eta}}{\sum_{j=1}^m \pi_j^* e^{\bar{S}^j/\eta}} \quad (4.24)$$

where π_i^* is model i ’s selection probability defined at its SCE.

Sketch of Proof. Since the experimentation probabilities are bounded away from 0 and 1, the chain is recurrent and ergodic. Tightness is therefore not an issue. What needs to be

done is to identify the limit. In principle, the same martingale method approach that was used to obtain the mean ODE in Proposition 4.2 can be applied. (See Appendix A). The main task is to verify that an appropriately defined ‘residual process’ is a martingale. \square

Proposition 4.11 asks what happens when $\tau \rightarrow \infty$ ‘faster’ than $\eta \rightarrow 0$. It’s also useful to ask what happens when $\eta \rightarrow 0$ for a given (large) τ . It’s clear from equation (4.24) that this limit is degenerate.

Proposition 4.12. *As $\eta \rightarrow 0$ the model distribution collapses onto the model with the largest LD rate function.*

This means that in the limit the decision maker will come to use one of the models ‘almost always’. The dominant model will be the model with the largest LD rate function. This model survives specification testing longer than any other model. Interestingly, the dominant model may not be the best fitting model. Of course, all else equal, poorly fitting models will have smaller rate functions and will not endure specification testing for long. A large residual variance generates a lot of ‘noise’ around the SCE, and therefore, makes it easier to escape. However, the *precision* of a model’s estimates also matters. Precise estimates are less liable to wander from their SCE values. Hence, overfitted models can escape just as quickly as underfitted models.

The reader may have noticed that we have not paid much attention to the details of randomization. Propositions 4.11 and 4.12 show why. It turns out that our LD approach embodies a form of robustness with respect to the details of experimentation. All that matters is that each model’s chances of being selected remain strictly bounded between 0 and 1.

Corollary 4.13. *As long as the experimentation probabilities, π_t^i , remain strictly bounded between 0 and 1, the identity of the dominant SCE is independent of the details of randomization.*

Proof. Follows directly from equation (4.24). \square

4.4.1. *An Information-Theoretic Interpretation.* We have defined a validated self-confirming equilibrium as an outcome generated by a model which survives specification testing longer than any other model in \mathcal{M} . We have identified this model as the model with the maximum large deviations rate function, defined at its own self-confirming equilibrium. To readers familiar with information theory and statistics, this may at first sight appear to be a puzzling result. From Sanov’s Theorem we know rate functions are connected to relative entropies, and then, from either Stein’s lemma (classical) or Chernoff bounds (Bayesian), we know that relative entropies are connected to detection error probabilities. In particular, larger relative entropies should make it easier to detect discrepancies between a model and the true DGP. That is, larger relative entropies reduce the probabilities of Type I and Type II errors. Why then are models with large rate functions *more* durable?

This apparent contradiction illustrates a key difference between model validation with exogenous data and model validation with endogenous data. With endogenous data, each model has the capacity to mimic the true DGP. In this case, rejecting a model constitutes a Type I error, and as usual, a larger rate function implies a smaller Type I error probability (or more precisely, it increases the rate at which it converges to zero).

4.4.2. *An Informal Example.* Suppose $\dim(\mathcal{M}) = 3$, i.e., there are 3 possible models. Let S_i^* be the large deviations rate function for model- i , evaluated at its unique stable SCE. The combination of constant gain learning, specification testing, and random model selection induces an approximating 3-state ergodic Markov chain across models. Model switches are triggered by escapes from each model's SCE. As $\eta \rightarrow 0$, these escape probabilities are of order $e^{-S_i^*/\eta}$. Model selection dynamics can therefore be approximated by the 3-state transition matrix, $\bar{P} = I + \eta Q^\eta$, where Q^η is the generator

$$Q^\eta = \begin{pmatrix} -(\pi_2^* + \pi_3^*)e^{-S_1^*/\eta} & \pi_2^*e^{-S_1^*/\eta} & \pi_3^*e^{-S_1^*/\eta} \\ \pi_1^*e^{-S_2^*/\eta} & -(\pi_1^* + \pi_3^*)e^{-S_2^*/\eta} & \pi_3^*e^{-S_2^*/\eta} \\ \pi_1^*e^{-S_3^*/\eta} & \pi_2^*e^{-S_3^*/\eta} & -(\pi_1^* + \pi_2^*)e^{-S_3^*/\eta} \end{pmatrix} \quad (4.25)$$

and where $\pi_i^* \in (0, 1)$ are parameters determining which model is more likely to be selected following a given model rejection.

The stationary distribution is as follows,

$$\begin{aligned} \bar{p}_1 &= \Delta^{-1} a_1 e^{-(S_2^* + S_3^*)/\eta} \\ \bar{p}_2 &= \Delta^{-1} a_2 e^{-(S_1^* + S_3^*)/\eta} \\ \bar{p}_3 &= \Delta^{-1} a_3 e^{-(S_1^* + S_2^*)/\eta} \end{aligned}$$

where

$$\Delta = a_1 e^{-(S_2^* + S_3^*)/\eta} + a_2 e^{-(S_1^* + S_3^*)/\eta} + a_3 e^{-(S_1^* + S_2^*)/\eta}$$

and where a_i are constants that are independent of η . Therefore,

$$\begin{aligned} \frac{\bar{p}_2}{\bar{p}_1} &\propto e^{-(S_1^* - S_2^*)/\eta} \\ \frac{\bar{p}_3}{\bar{p}_1} &\propto e^{-(S_1^* - S_3^*)/\eta} \end{aligned}$$

Suppose Model 1 is dominant, so that $S_1^* > S_2^*$ and $S_1^* > S_3^*$. Then notice that as $\eta \rightarrow 0$, Model 1 is used almost always, and this conclusion does not depend on the experimentation probabilities. This independence derives from the fact that once a model starts to be used, its coefficient estimates converge to their self-confirming equilibrium values.

4.5. A Return to the Phillips Curve. At this point, it is of some interest to revisit the Phillips Curve example in Section 2 from the perspective of our theory of model validation. We saw that testing hypotheses within the context of a single model would not produce good outcomes. The Bank continues to rediscover a statistical Phillips Curve, and continues to let inflation slip out of control. We then saw how good theory, in the form of a distinction between expected and unexpected inflation, could produce good outcomes. Suppose, however, the bright young theorist from Section 2 never does get hired, but instead the Bank follows the kind of model validation procedure we've just outlined. As Sims (1982) argued, could the Bank learn to adopt the right policy even without a priori knowledge of the DGP?

Suppose the Bank begins with a model class comprised of *two* elements: (1) a statistical Phillips curve, as in Sargent (1999), and (2) a vertical Phillips curve. Our theory predicts

that the dominant model will be the model with biggest large deviations rate function. In this particular case, this is an easy comparison to make. We know from Cho, Williams, and Sargent (2002) that the static Phillips Curve rate function is approximately $\bar{S}^* = .0005$. The mean escape time is approximately $\exp[\bar{S}^*/\eta]$ continuous time units, or $\eta^{-1} \exp[\bar{S}^*/\eta]$ discrete time units.¹⁸ Hence, when $\eta = .01$ we should expect to observe escapes every 105 periods.¹⁹ If the LM test is calibrated to reject only during escapes, then the static Phillips Curve would be expected to last about 100 periods. The vertical Phillips Curve case is especially simple. Since the sequence of coefficient estimates becomes Gaussian, the rate function is well known to be $\bar{S}^*(x) = .5(x - \bar{u})^2 / (\sigma_1^2 + \sigma_2^2)$, where σ_1^2 and σ_2^2 are the variances of the shocks to the Phillips Curve and the inflation target, and \bar{u} is the natural rate of unemployment. Note that in this linear setting, the rate function is symmetric, and escapes are equally likely to occur in either direction. To maintain comparability with the static Phillips Curve we need to calibrate the boundary point, x , so that model rejections occur only during escapes, and with approximately equal statistical evidence. From Figure 7, rejections of the static Phillips Curve occur when the LM test reaches levels of approximately 16. Since in the case of a vertical Phillips Curve, the LM test essentially becomes a recursive F-test, or a squared t -statistic, this suggests a mean escape time of approximately $\eta^{-1} \exp[8]$ discrete time units; that is, about once every 300,000 periods! Clearly, the vertical Phillips Curve would dominate, and for all practical purposes the Bank would stick to a low inflation policy forever.

5. CONCLUDING REMARKS

This paper has attempted to model macroeconomic policymakers as econometricians. We've done this by combining recent work in both macroeconomics and econometrics. From macroeconomics, we've borrowed from the work of Sargent (1999) and Evans and Honkapohja (2001) on boundedly rational learning dynamics. From econometrics, we've borrowed from recent work on the analysis of misspecified models (Vuong (1989), White (1994), and Burnham and Anderson (2002)). As it turns out, this produces a rather difficult, and as yet unconsummated, marriage.

From a macroeconomic standpoint, it is difficult because we abandon the Rational Expectations Hypothesis, thereby putting ourselves into the 'wilderness of bounded rationality'. We do this not because we like to analyze difficult and ill-posed problems, but simply because of the casual observation that, as econometricians, macroeconomic policymakers do not spend their time refining estimates of a known model, but instead spend most of their time searching for new and better models. Of course, it is not *necessary* to abandon Rational Expectations and traditional Bayesian decision theory when confronting model uncertainty.²⁰ However, we think there are good reasons to explore alternative approaches.²¹

¹⁸Warning: For Gaussian shocks this formula provides only an upper bound. It becomes an equality only when shocks are bounded.

¹⁹Another warning: The distribution of escape times is not symmetric. It is exponential, with a long right tail. Hence, the median escape time will be less than this.

²⁰See, e.g., Brock, Durlauf, and West (2007) for an application of Bayesian model averaging to macroeconomic policy.

²¹See Sargent (1993), Hansen and Sargent (2007b), Kreps (1998), and Bray and Kreps (1987).

The marriage between macroeconomics and econometrics is difficult from an econometric standpoint because, presumably, policymakers have some influence over the data-generating processes they are attempting to learn about. The econometric analysis of misspecified models with endogenously generated data is truly uncharted territory.

We make progress on this problem by relating it to a problem that *is* relatively well understood, namely, the dynamics of constant gain recursive learning algorithms. We describe the sense in which the dynamics generated by a process of testing and model revision can be approximated by the dynamics generated by recursive learning models. This is a useful connection to make, because it enables us to apply the results of Williams (2001) and Cho, Williams, and Sargent (2002) on escape dynamics to help us understand a wide range of Markov-switching macroeconomic dynamics. Looking at it from the other side, a second payoff from making this connection is that it provides a more secure behavioral foundation for recursive learning models.

Although we feel this paper takes a significant step forward in understanding the interplay between macroeconomics and econometrics, there are certainly many loose ends and unexplored avenues remaining. One possibility is to consider alternative specification tests. Here we focused on LM tests. However, there are many possibilities, depending on what sort of potential misspecification is of most concern. Another possible extension would be to consider alternative model classes. The analysis here was confined to simple atheoretical linear regression models. However, it is becoming increasingly common to base policy discussions on more sophisticated structural models. The second step of our validation process, which involves comparing alternative models in the event the current reference model is rejected, might benefit from recent advances in Bayesian model comparison (Schorfheide (2000)). Third, like most work on bounded rationality, our analysis rests on an uneasy tension between positive and normative ingredients. For the most part, we have struck the balance on the positive side, and adopted assumptions that are meant to be more descriptive than prescriptive. The one major exception here was our assumption about model revision. For the sake of analytical convenience, we assumed the decision maker was quite astute at recognizing the counterfactual dilemmas posed by learning in self-referential environments. However, assuming that models are only fit to data that were generated while they were used is a bit of a stretch, to put it mildly. It would be useful to relax this assumption in some way. Finally, perhaps the most interesting and important extension would be to allow the agent to entertain doubts about the entire model class itself. The work of Hansen and Sargent (2007a) on robust filtering of discrete hidden states offers one route toward such an extension. Another possibility is to take advantage of recent work on calibrated learning in repeated games. We are currently pursuing this in Cho and Kasa (2008).

APPENDIX A. PROOF OF PROPOSITION 4.2

There are two key steps to any weak convergence argument: (1) Establish tightness, and (2) Identify the limit. Tightness delivers compactness (in the space of right-continuous functions with left-hand limits, endowed with the Skorohod topology (Prohorov's Theorem)), which then guarantees existence of a weakly convergent subsequence. Proving tightness can be challenging. However, given our level of generality, we simply assume it, since the details of any proof are unavoidably case specific. One can always guarantee it by resort to a projection facility.

To identify the limit, we can employ the powerful martingale method (Kushner and Yin (1997)). The logic of this method is as follows. Let β_i^* be the SCE of model- i . Let β_n^i denote the discrete-time sequence of coefficient estimates, and $\beta^i(t)$ be its continuous time interpolation. If we then define $\tilde{\beta}_n^i = \beta_n^i - \beta_i^*$ and $\tilde{\beta}_\eta^i(t) = \beta^i(t) - \beta_i^*$ as the deviations of the estimates from the SCE, we can show

$$\lim_{\eta \rightarrow 0} \tilde{\beta}_\eta^i(t) \Rightarrow \bar{\beta}^i(t)$$

where $\bar{\beta}^i(t)$ is the solution to

$$\bar{\beta}^i(t) = \bar{\beta}^i(0) + \int_0^t h^i(\bar{\beta}^i(s)) ds$$

and where $h^i(\cdot)$ defines the mean dynamics of model- i , obtained by averaging out the data with respect to its (model specific) invariant distribution at the given value of β^i . To establish this, we need a key technical result, viz., that any continuous-time martingale with Lipschitz continuous paths (w.p.1) is a constant (again, w.p.1). (See Kushner and Yin (1997) for a proof). To exploit this result, define the following process

$$W^\eta(t) = \tilde{\beta}_\eta^i(t) - \beta_0^i - \eta \sum_{n=0}^{t/\eta-1} h^i(\beta_n^i)$$

and let

$$W(t) = \bar{\beta}^i(t) - \bar{\beta}^i(0) - \int_0^t h^i(\bar{\beta}^i(s)) ds$$

be its limit, which exists by Prohorov's Theorem. If we can prove $W(t)$ is a martingale then we're done, since $W(0) = 0$. We shall not provide all the details, since they are fairly tedious, and routine in any case.

APPENDIX B. PROOF OF LEMMA 4.8

We must show that $\forall \epsilon > 0$ there exists a $\delta > 0$ such that $\|\beta_1 - \beta_2\| < \delta$ implies $\|\theta_1 - \theta_2\| < \epsilon$, where $\|\cdot\|$ denotes the sup norm on $C[0, T]$. Suppose $|\mathcal{Z}(\beta_{1,n}) - \mathcal{Z}(\beta_{2,n})| < M|\beta_{1,n} - \beta_{2,n}|$. By definition,

$$\begin{aligned} |\theta_1(t) - \theta_2(t)| &= \left| \int_0^t [\mathcal{Z}(\beta_1(s)) - \theta_1(s)] ds - \int_0^t [\mathcal{Z}(\beta_2(s)) - \theta_2(s)] ds \right| \\ &\leq \int_0^t |\mathcal{Z}(\beta_1(s)) - \mathcal{Z}(\beta_2(s))| ds + \int_0^t |\theta_2(s) - \theta_1(s)| ds \\ &\leq M\delta t + \int_0^t |\theta_2(s) - \theta_1(s)| ds \end{aligned}$$

Applying Gronwall's lemma to the last inequality gives

$$|\theta_1(t) - \theta_2(t)| \leq M\delta t + M\delta e^t \int_0^t s e^{-s} ds$$

and the result follows.

REFERENCES

- BENVENISTE, A., M. METIVIER, AND P. PRIOURET (1990): *Adaptive Algorithms and Stochastic Approximations*. Springer-Verlag, Berlin.
- BRANCH, W. A., AND G. W. EVANS (2006): “Intrinsic Heterogeneity in Expectation Formation,” *Journal of Economic Theory*, 127, 264–295.
- (2007): “Model Uncertainty and Endogenous Volatility,” *Review of Economic Dynamics*, 10, 207–37.
- BRAY, M., AND N. SAVIN (1986): “Rational Expectations Equilibria, Learning, and Model Specification,” *Econometrica*, 54, 1129–60.
- BRAY, M. M. (1982): “Learning, Estimation and the Stability of Rational Expectations,” *Journal of Economic Theory*, 26, 318–339.
- 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.
- BREUSCH, T. S., AND A. R. PAGAN (1980): “The Lagrange Multiplier Test and its Applications to Model Specification in Econometrics,” *Review of Economic Studies*, 47, 239–253.
- BROCK, W. A., S. N. DURLAUF, AND K. D. WEST (2007): “Model Uncertainty and Policy Evaluation: Some Theory and Empirics,” *Journal of Econometrics*, 136, 629–64.
- BROCK, W. A., AND C. H. HOMMES (1997): “A Rational Route to Randomness,” *Econometrica*, 65, 1059–1095.
- BULLARD, J., AND K. MITRA (2002): “Learning About Monetary Policy Rules,” *Journal of Monetary Economics*, 49, 1105–29.
- BURNHAM, K., AND D. ANDERSON (2002): *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*. Springer, 2nd edn.
- CHEN, X., AND H. WHITE (1998): “Nonparametric Adaptive Learning with Feedback,” *Journal of Economic Theory*, 81.
- CHO, I.-K., AND K. KASA (2008): “Calibrated Learning and Robust Model Validation,” University of Illinois.
- CHO, I.-K., N. WILLIAMS, AND T. J. SARGENT (2002): “Escaping Nash Inflation,” *Review of Economic Studies*, 69, 1–40.
- CHU, J., M. STINCHCOMBE, AND H. WHITE (1996): “Monitoring Structural Change,” *Econometrica*, 64, 1045–1065.
- COGLEY, T., R. COLACITO, L. P. HANSEN, AND T. J. SARGENT (2008): “Robustness and U.S. Monetary Policy Experimentation,” New York University.
- COGLEY, T., AND T. J. SARGENT (2005): “The Conquest of U.S. Inflation: Learning and Robustness to Model Uncertainty,” *Review of Economic Dynamics*, 8, 528–63.
- DAWID, A. P. (1984): “Statistical Theory: The Prequential Approach,” *Journal of the Royal Statistical Society, Series A*, 147, 278–292.
- DEMBO, A., AND O. ZEITOUNI (1998): *Large Deviations Techniques and Applications*. Springer-Verlag, New York, 2nd edn.
- DESHAYES, J., AND D. PICARD (1982): “Tests of Disorder of Regression: Asymptotic Comparison,” *Theory of Probability and its Applications*, 27, 100–15.
- ENGLE, R. F. (1984): “Wald, Likelihood Ratio, and Lagrange Multiplier Tests in Econometrics,” in *Handbook of Econometrics, Volume II*. Elsevier Science Publishers.
- EVANS, G., AND S. HONKAPOHJA (2003): “Expectations and the Stability Problem for Optimal Monetary Policy,” *Review of Economic Studies*, 70, 807–24.
- EVANS, G. W., AND S. HONKAPOHJA (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press.
- FOSTER, D. P., AND R. V. VOHRA (1998): “Asymptotic Calibration,” *Biometrika*, 85, 379–390.
- FOSTER, D. P., AND H. P. YOUNG (2003): “Learning, Hypothesis Testing and Nash Equilibrium,” *Games and Economic Behavior*, 45, 73–96.
- FUDENBERG, D., AND D. K. LEVINE (2008): “Self-Confirming Equilibrium and the Lucas Critique,” forthcoming, *Journal of Economic Theory*.

- GILBOA, I., A. W. POSTLEWAITE, AND D. SCHMEIDLER (2008): "Probability and Uncertainty in Economic Modeling," *Journal of Economic Perspectives*, 22, 173–188.
- HANSEN, B. E. (1992): "Testing for Parameter Instability in Linear Models," *Journal of Policy Modeling*, 14, 517–533.
- HANSEN, L. P., AND T. J. SARGENT (1993): "Seasonality and Approximation Errors in Rational Expectations Models," *Journal of Econometrics*, 55, 21–55.
- (2007a): "Recursive Robust Estimation and Control Without Commitment," *Journal of Economic Theory*, 136, 1–27.
- (2007b): *Robustness*. Princeton University Press.
- HANSEN, M. H., AND B. YU (2001): "Model Selection and the Principle of Minimum Description Length," *Journal of the American Statistical Association*, 96, 746–774.
- KANDORI, M., G. MAILATH, AND R. ROB (1993): "Learning, Mutation and Long Run Equilibria in Games," *Econometrica*, 61, 27–56.
- KOCHERLAKOTA, N. R. (2007): "Model Fit and Model Selection," *Federal Reserve Bank of St. Louis Review*, 89, 349–60.
- KREPS, D. M. (1998): "Anticipated Utility and Dynamic Choice," in *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983-1997*. Cambridge University Press.
- 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.
- RIVERS, D., AND Q. VUONG (2002): "Model Selection Tests for Nonlinear Dynamic Models," *Econometrics Journal*, 5, 1–39.
- SARGENT, T. J. (1993): *Bounded Rationality in Macroeconomics*. Clarendon Press.
- (1999): *The Conquest of American Inflation*. Princeton University Press.
- SARGENT, T. J., AND N. WILLIAMS (2005): "Impacts of Priors on Convergence and Escapes from Nash Inflation," *Review of Economic Dynamics*, 8, 360–391.
- SCHORFHEIDE, F. (2000): "A Loss-Function Based Evaluation of DSGE Models," *Journal of Applied Econometrics*, 15, 645–70.
- SIMS, C. A. (1982): "Policy Analysis with Econometric Models," *Brookings Papers on Economic Activity*, 1:1982, 107–164.
- SIN, C.-Y., AND H. WHITE (1996): "Information Criteria for Selecting Possibly Misspecified Parametric Models," *Journal of Econometrics*, 71, 207–225.
- VUONG, Q. H. (1989): "Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses," *Econometrica*, 57(2), 307–333.
- WHITE, H. (1982): "Maximum Likelihood Estimation of Misspecified Models," *Econometrica*, 50, 1–25.
- (1994): *Estimation, Inference and Specification Analysis*. Cambridge University Press.
- WILLIAMS, N. (2001): "Escape Dynamics in Learning Models," Ph.D. thesis, University of Chicago.