

LEARNING AND MODEL VALIDATION

IN-KOO CHO AND KENNETH KASA

ABSTRACT. This paper addresses an important weakness of the macroeconomic learning literature by allowing agents to test the specification of their models. In contrast to existing macroeconomic learning models, we assume agents are aware of potential model misspecification, and try to detect it, in real-time, using an econometric specification test. If a model fails the test, a new model is formulated. If a model passes the test, it is used to implement a policy based on the provisional assumption that the current model is correctly specified, and will not change in the future. At the same time, policy decisions feedback to influence the model that agents are trying to learn about.

Although this testing and model validation process is an accurate description of most macroeconomic policy problems, 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

Macroeconomic policymaking confronts two serious challenges. First, policymakers do not know the model. That is, they do not know the economy’s actual data-generating process (DGP). Second, their beliefs about the model typically feedback to influence the DGP. That is, policy operates within a ‘self-referential’ system.

Considered separately, each of these difficulties can be analyzed with standard methods. In the absence of feedback, model uncertainty can be addressed using a variety of approaches. For example, during the past two decades classical econometricians have made substantial progress on testing and comparing misspecified models.¹ At the same time, there has been an explosion of recent work on Bayesian macroeconometrics.² For Bayesians, there is no essential difference between model uncertainty and parameter uncertainty. Specification of a prior over models essentially embeds model uncertainty within

Date: August, 2009.

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.

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

²An and Schorfheide (2007) contains a summary.

a larger space consisting of a single ‘hypermodel’. Unfortunately, genuine model uncertainty is infinite dimensional, and specifying priors over infinite dimensional objects is not easy. Recent work on robust control opens the door to an analysis of model uncertainty in situations where it may be difficult to specify a prior over models (Hansen and Sargent (2007a)). However, none of these approaches permits feedback from beliefs about models to the actual DGP. The first two adopt the perspective of an outside observer, while the latter typically assumes the DGP is a fixed reference model.

Of course, feedback is nothing new to economists. It lies at the heart of the Lucas Critique, and macroeconometricians have developed methods for coping with it (Lucas and Sargent (1981)). However, these methods presume that all agents know the model, including the policymaker. Although the macroeconomic learning literature is designed to relax this assumption, to date it either assumes a priori knowledge of the correct model (at least up to a small handful of unknown parameters), or it assumes that any model misspecification goes unquestioned and undetected.³

Our paper attempts to confront both challenges simultaneously. Its goal is to understand how feedback between hypothesis testing and the DGP influences 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.⁴ However, it is not at all obvious whether this is indeed true when the data are endogenous. 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 (SCE). As in Fudenberg and Levine (2008), we are interested in studying the role of experimentation in selecting favorable SCE.

We assume policymaking takes place via the following four step process: (1) A policymaker is endowed with a fixed *model class*, consisting of a collection of potentially misspecified linear models, each with a finite number of unknown parameters, (2) Each period the policymaker 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 randomly selected from the model class. This combined process of estimation, testing, and selection will be called *model validation*.

The last step of this algorithm is crucial. It implies the policymaker *deliberately* injects randomness into his model choice. This randomization does not reflect capriciousness or computational errors, but instead reflects a strategic response to model uncertainty. In particular, it prevents ‘nature’ from exploiting the policymaker’s ignorance of the DGP, and in this way delivers a form of ‘robustness’ (Foster and Vohra (1998)). It can also be interpreted as a form of experimentation. Our main conclusions do not depend on the details of this experimentation process.

³See Evans and Honkapohja (2001) for a survey of the macro learning literature. Sargent (1999) contains an interesting example of learning and macroeconomic policy with a misspecified model.

⁴See, in particular, pg. 117 of Sims (1982).

Note that in our 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 (2009). The dynamics of recursive model comparison have some subtle features not found in validation dynamics, 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.

In conventional learning models (e.g., Marcet and Sargent (1989)), one can analyze the asymptotic properties of the learning dynamics by approximating the evolution of a model's coefficients by the trajectory of an ordinary differential equation (ODE). This is due to the fact that coefficient estimates evolve on a slower time-scale than the data. A stationary solution of this ODE defines a self-confirming equilibrium, and its stability provides a natural equilibrium selection criterion.

With model validation, however, models are occasionally replaced by different models. As a result, the relevant state is defined not only by coefficient estimates, but also by the realization of a discrete model indicator. In principle, the evolution of this enlarged state space is considerably more complex. In general, it cannot be approximated by an ODE. Fortunately, by appropriately calibrating the testing threshold, we can exploit the same sort of time-scale separation as in conventional learning models. This is possible due to the endogenous nature of the DGP. Specifically, in the neighborhood of a stable self-confirming equilibrium, the probability of a model rejection will be very small. Hence, models survive specification testing for a long time, where 'long' is defined relative to the rate of coefficient updating. This allows the evolution of each model's coefficients to be approximated separately, using the trajectory of a conventional ODE.

More precisely, by exploiting the large deviation properties of each model's self-confirming equilibrium, we show that as the coefficient update gain decreases, a model's coefficients remain close to its self-confirming equilibrium for an *exponentially* long period of time. As long as the policymaker tries each model with positive probability following model rejections, a self-confirming equilibrium must be visited within a finite amount of time. Hence, as the gain converges to zero, the policymaker spends an increasing fraction of time using a *single* model. Large deviations theory allows us to identify this model. This kind of result bears some resemblance to the evolutionary game theory literature (Kandori, Mailath, and Rob (1993)).

Large deviation methods also deliver an interesting and useful interpretation of this limiting model. Specifically, we show that it is the model possessing the largest so-called 'rate function'. A key result in the theory of large deviations (i.e., Sanov's theorem) links this rate function to relative entropy and the Kullback-Leibler Information Criterion (KLIC). This is interesting, since the KLIC is a pervasive concept in the econometrics literature on model testing and selection. The relative entropy that is being captured by each model's rate function can be interpreted as the KLIC distance between the probability distribution associated with its SCE and the distribution associated with the closest model that triggers a rejection or escape. This extends the results of White (1982) in a natural way to the case of endogenous data.

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, asking whether agents could use standard diagnostics, like Chow tests and Durbin-Watson statistics, to detect the parameter variation 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.

The remainder of the paper is organized as follows. Section 2 provides an example 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 applies our results to Sargent's (1999) Phillips Curve model. Section 6 contains a few concluding remarks, and an appendix contains proofs of some technical results.

2. AN EXAMPLE

To set the stage, we first present an example designed to highlight why we think it is important to incorporate specification testing and model validation procedures into the macroeconomic learning literature. The example also illustrates the difference between hypothesis testing and recursive model selection.

2.1. 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 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 robust than others.

Suppose the true price data generating process is given by

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

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 and mutually independent. This model has the following rational

expectations equilibrium,

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

Following Evans and Honkapohja (2001), suppose the agent entertains two different models, each obtained by including only one element of $\{w_1, w_2\}$ on the right hand side of the equation.⁵ We thus endow the agent with a model class $\mathcal{M} = \{\mathcal{M}_1, \mathcal{M}_2\}$, where

$$\mathcal{M}_1 : \quad p_n = \beta_1 w_{1,n-1} + v_{1,t} \quad (2.2)$$

$$\mathcal{M}_2 : \quad p_n = \beta_2 w_{2,n-1} + v_{2,t} \quad (2.3)$$

The agent believes the disturbance process in each model is i.i.d, and orthogonal to the included regressor. Hence, each model is estimated using recursive least squares,

$$\hat{\beta}_{1,n} = \hat{\beta}_{1,n-1} + \eta R_{1,n-1}^{-1} w_{1,n-1} (p_n - \hat{\beta}_{1,n-1} w_{1,n-1})$$

$$\hat{\beta}_{2,n} = \hat{\beta}_{2,n-1} + \eta R_{2,n-1}^{-1} w_{2,n-1} (p_n - \hat{\beta}_{2,n-1} w_{2,n-1})$$

$$\hat{R}_{1,n} = \hat{R}_{1,n-1} + \eta (w_{1,n}^2 - \hat{R}_{1,n-1})$$

$$\hat{R}_{2,n} = \hat{R}_{2,n-1} + \eta (w_{2,n}^2 - \hat{R}_{2,n-1})$$

where the constant gain parameter, $0 < \eta < 1$, serves to discount old data. This reflects the agent's doubts about the stationarity of the environment, which turn out to be self-confirming.

Note that in this model the agent only takes two actions: (1) He selects a model at the beginning of each period, and (2) He uses it to construct a forecast of next period's price. These forecasts feedback to influence the actual price process as follows:

$$p_n = \alpha (s_{1,n-1} \hat{\beta}_{1,n-1} w_{1,n-1} + (1 - s_{1,n-1}) \hat{\beta}_{2,n-1} w_{2,n-1}) + \beta_1 w_{1,n-1} + \beta_2 w_{2,n-1} + \varepsilon_n$$

where $s_{1,n}$ is a model indicator, with $s_{1,n} = 1$ when \mathcal{M}_1 is used in period- n , and zero otherwise. In the next section we develop a more general framework, in which models are used to formulate policy actions, and these policy actions feedback to influence the true data-generating process.

Despite its simplicity, this model does raise an important question; namely, to what extent are policymakers aware of their own influence over the data-generating process? If they are aware of this influence, they face 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 misleading inferences about the prospects of a given model. One way to account for this would be to only update models while they are in use. The above algorithm strikes a more realistic balance. Although models are continuously updated, even when they are not being used, the fact that data are discounted means that data generated by previous models gets downweighted.

It turns out that inferences about long-run model use do not depend on how unused models are updated. In the long-run, one model dominates, and so its coefficients are

⁵To facilitate comparison to Branch and Evans (2007), suppose the agent cannot access a correctly specified model.

estimated with an ever larger fraction of data that were self-generated. In a sense, the fact that a model is updated using data from another model is evidence that this model is dominated. What's important is the the combination of experimentation and constant gain parameter updating. Experimentation guarantees the self-confirming equilibria of all models get a chance to prove themselves. Without experimentation, the agent may overlook models simply because they perform poorly when fit to data generated by other models. At the same time, constant gain updating implies that all models will eventually be rejected. This prevents the agent from getting trapped by inferior self-confirming equilibria.

While the set-up here is similar Branch and Evans (2006), their goal is quite different. 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 continuously challenges the existing model. Hypothesis testing generates model switches.

Suppose that $\alpha < 1$. Branch and Evans (2006) show that when model \mathcal{M}_1 is used exclusively it has a stable self-confirming equilibrium,⁶

$$p_n = \frac{1}{1-\alpha}(\beta_1 + \Omega_{11}^{-1}\Omega_{12}\beta_2)w_{1,n-1} + \varepsilon_n \quad (2.4)$$

while model \mathcal{M}_2 has its own stable self-confirming equilibrium

$$p_n = \frac{1}{1-\alpha}(\beta_2 + \Omega_{22}^{-1}\Omega_{21}\beta_1)w_{2,n-1} + \varepsilon_n \quad (2.5)$$

where Ω_{ij} are the elements of the second moment matrix, $\Omega = E(ww')$. Because the self-confirming equilibria (2.4) and (2.5) are obtained by restricting the agent's beliefs to a misspecified model, Evans and Honkapohja (2001) refer to them as Restricted Perception Equilibria (RPE). We take this idea a step further. In a sense we conduct a meta-stability analysis, and ask which of several stable RPE is likely to dominate in the long-run when agents are allowed to test the specification of their models.

In the context of this particular example, we can now describe the details of our model validation process

- (1) The agent enters period- n with a given model $\ell \in \{1, 2\}$, with coefficients that have been updated based upon all observations up through period $n - 1$.
- (2) After observing p_n , the agent computes each model's squared forecast error, $\xi_{\ell,n} = (p_n - \hat{\beta}_{\ell,n-1}w_{\ell,n-1})^2$, and then updates his estimate of each model's mean-squared forecast error,

$$\hat{\sigma}_{\ell,n}^2 = \hat{\sigma}_{\ell,n-1}^2 + \eta(\xi_{\ell,n} - \hat{\sigma}_{\ell,n-1}^2) \quad (2.6)$$

He continues to use model ℓ if and only if

$$\hat{\sigma}_{\ell,n}^2 \leq \theta$$

for a given threshold $\theta > 0$.

⁶Stable in the sense of E-stability (Evans and Honkapohja (2001)).

- (3) If model ℓ passes the test, the agent uses it to forecast next period's price. If model ℓ fails the test, the agent selects another model in \mathcal{M} . If this new model passes the test, it becomes the reference model, and is used to forecast next period's price. Otherwise the search continues until a successful model is discovered.
- (4) Move to period $n + 1$.

Note that in principle step (3) could continue indefinitely, if there is no model in \mathcal{M} that can pass the test. In this particular example we avoid this trivial outcome by calibrating the test threshold so as to guarantee that each model passes the test while in the neighborhood of its own self-confirming equilibrium. In the next section we outline a more sophisticated statistical test, based on Lagrange Multiplier testing principles, which due to the endogeneity of the data, guarantees that each model will pass the test at its own self-confirming equilibrium.⁷

Figure 1 shows what happens when an agent tests his model each period, rejecting it if its mean-squared error exceeds a fixed threshold.⁸ If a model is rejected, the agent selects randomly between them with weights determined by a logit function. 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.⁹




FIGURE 1. Model Validation in a Misspecified Cobweb Model

Not surprisingly, the agent uses model 1 most of the time. In fact, from the model weights in the lower panel, we can see that model 1 is sometimes used even in periods when model 2 has a lower estimated mean-squared error. This illustrates an important difference between specification testing and recursive model selection.

Although model 1 is used most of the time, notice that occasionally a prolonged excursion to Model 2 takes place. These excursions represent the model's *escape dynamics*.

⁷While more sophisticated than a simple comparison of mean-squared errors with a fixed threshold, a Lagrange Multiplier test is not sophisticated enough to detect some subtle forms of misspecification. This will become clear in our discussion of Sargent's example in section 5.

⁸Leave aside for the moment the naivete associated with repeatedly testing against a fixed threshold. This will be addressed in the following examples, and makes little difference to the point being made here.

⁹Some more details: (1) w_1 and w_2 are assumed independent with $\sigma_{w_1}^2 = .4$ and $\sigma_{w_2}^2 = .3$, (2) $\sigma_\varepsilon^2 = .25$, (3) $\beta_1 = \beta_2 = 1.0$, (4) $\alpha = 0.4$, (5) The update gain is set to $\eta = .03$, (6) The testing threshold is set to $\tau = 1.0$, and (7) The 'choice intensity' parameter of the logit function is set to 15.0.

They can be analytically characterized using the tools of large deviations theory. The RPE of each model has a rate function attached to it. These rate functions summarize how unlikely it is to escape from each of the RPE. In the limit, as the update gain gets smaller, Model 1 becomes dominant, and is used an increasingly large fraction of the time. Accordingly, we refer to the RPE of model 1 as a *robust* self-confirming equilibrium. We show that robust self-confirming equilibria possess relatively large rate functions.

3. A GENERAL FRAMEWORK

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

We study a policymaker 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 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 (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.¹⁰ 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.¹¹

3.1. Objective Function. Consider a policymaker 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,

$$\mathbf{E}_n \sum_{j=0}^{\infty} \delta^j \{x'_{n+j} Q x_{n+j} + u'_{n+j} R u_{n+j}\} \quad (3.7)$$

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 policymaker. 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 policymaker 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. Forward-looking equations can be accommodated by adding Lagrange multipliers to the state. Our maintained assumption is that the policymaker believes there to be a

¹⁰Brock, Durlauf, and West (2007) make progress in applying Bayesian methods to macroeconomic policy design.

¹¹Interestingly, 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.

single ‘best’ model within \mathcal{M} , 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 policymaker 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 policymaker to reconsider his original specification of \mathcal{M} .¹²

We assume the policymaker 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 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

$$x_{jn}^i = \alpha_j^i e_j^i y_{n-1} + \gamma^i e_{u,j}^i u_{n-1} + \varepsilon_{jn}^i \quad i = 1, 2, \dots, m \quad j = 1, 2, \dots, s \quad (3.8)$$

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

$$x_n = A_1(s_{n-1}, \beta_{n-1})x_{n-1} + B_1(s_{n-1}, \beta_{n-1})z_{n-1} + C_1(s_{n-1}, \beta_{n-1})u_n + D_1(s_{n-1}, \beta_{n-1})\epsilon_{1n} \quad (3.9)$$

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

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. However, all we need to know about the DGP is following:

Assumption 3.1. *For a fixed $(s, \beta) = (s_n, \beta_n)$, the joint distribution of (x_n, z_n) 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 require this assumption.

¹²Of course, this is only an issue if events seem to call for an expansion of the model class. By adjusting the weights, the policymaker can effectively discard (relatively) poorly performing models.

3.4. Model Updating. While in use, 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 3.8 as follows

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

where \mathbf{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 $\Phi_{n-1}^i = \text{diag}[\mathbf{E}_1^i \phi_{n-1}, \mathbf{E}_2^i \phi_{n-1}, \dots, \mathbf{E}_s^i \phi_{n-1}]$. 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

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

$$R_n^i = R_{n-1}^i + \eta (\Phi_{n-1}^i \Phi_{n-1}^{i'} - R_{n-1}^i) \quad (3.11)$$

where η is a constant gain parameter, assumed to be common across equations and models.

To facilitate the analysis, we assume that the set of feasible coefficients for each model satisfies some regularity conditions.

Assumption 3.2. *Let \mathcal{B}^i be the set of all feasible coefficients for model i . We assume that \mathcal{B}^i is compact and convex.*

The compactness assumption can be interpreted as a priori knowledge about the DGP, which rules out outrageously large coefficients. Convexity is mainly a technical assumption, designed to address the learning dynamics along the boundary of the parameter space.

3.5. Self-Confirming Equilibria. As in Marcet and Sargent (1989), the sample path of an individual model's coefficients can be approximated by the trajectory of an ordinary differential equation (ODE), conditional on the assumption that the data is generated by that model. We write each model's resulting ODE

$$\dot{\beta}^i = h^i(\beta^i) \quad (3.12)$$

where h^i is assumed to be Lipschitz continuous.

We can now state the following important definition:

Definition 3.3. *A self-confirming equilibrium (SCE) for model i consists of a fixed vector of regression coefficients, β_e^i , and a control feedback matrix, $F^i(\beta_e^i)$, which satisfy the following two conditions:*

- (a): *The coefficients are a stationary point of the mean ODE, $h^i(\beta_e^i) = 0$.*
- (b): *The control policy $u_n^i = -F^i(\beta_e^i)y_{n-1}$ solves the LQR in (3.7) for model i*

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

Assumption 3.4. *Each model $i = 1, 2, \dots, m$ has a unique stable self-confirming equilibrium. Let β_e^i be the self-confirming equilibrium of model i , and \mathcal{D}_e^i be the domain of attraction for β_e^i of (3.12). We assume that $\mathcal{D}_e^i = \mathcal{B}^i$, implying global stability of β_e^i . In addition, β_e^i is in the interior of \mathcal{B}^i .*

Uniqueness is not essential. If model i has multiple locally stable self-confirming equilibria, we can treat different self-confirming equilibria as different models, and invoke exactly the same logic. We only have to assume instead that \mathcal{B}^i is partitioned to the different domains of attraction to different self-confirming equilibria. This condition is mainly to avoid the presence of a stable orbit, where the decision maker will never learn the self-confirming equilibrium, even if he uses the same model indefinitely. In many economic applications of linear rational expectations models, global convergence is satisfied.

We regard the stability of a self-confirming equilibrium as a primary requirement of a sensible solution. Our task is to provide a selection criterion among locally stable equilibria by checking their resilience to specification testing.

Given the discrete time process $\{\hat{\beta}_n^i\}$ and the fixed gain parameter $\eta > 0$, we can construct a continuous time process $\hat{\beta}_\eta^i(\tau)$ by linear interpolation. For $\tau \geq 0$, define

$$k = \min\{n | n\eta > \tau\}.$$

Then, define

$$\hat{\beta}_\eta^i(\tau) = \frac{k\eta - \tau}{\eta} \hat{\beta}_{k-1}^i + \frac{\tau - (k-1)\eta}{\eta} \hat{\beta}_k^i.$$

We follow the same convention to describe other continuous time processes. We refer to $n = 1, 2, \dots$ as periods or iteration numbers, while we refer to τ as fictitious time, or simply time.

3.6. Model Validation. We have been purposely vague about the details of specification testing, 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. 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.¹³ Essentially then, what is being tested is the significance of the algorithm's update term.

¹³Breusch and Pagan (1980) noted this as well, and argued that it could be exploited to simplify the computation of conventional (nonrecursive) LM tests.

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 2 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.¹⁴




FIGURE 2. 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 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. Fortunately, our main conclusions do not depend on the details of the threshold. What’s important is that the *same* threshold apply to all models.

Before proceeding, it will be useful to make explicit the connection between our model update equations and a traditional scoring algorithm by rewriting equation (3.10) as follows:

$$\hat{\beta}_n^i = \hat{\beta}_{n-1}^i + \eta(R_{n-1}^i)^{-1} \Lambda_n^i \quad (3.13)$$

¹⁴Some 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 2 then reports the recursively estimated statistic, $\mathcal{L}_n = \mathcal{L}_{n-1} + \eta[(x_n' \xi_n) R_n^{-1} (x_n \xi_n) / \hat{\sigma}_n^2 - \mathcal{L}_{n-1}]$.

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 .

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 main steps, where the first step consists of a misspecification test, and the second step specifies what happens in the event of a rejection.

We assume that once a given model is rejected, a new model, along with its coefficients, is randomly selected from the model class. Let π^i be the probability of selecting model i , and $F^i(\beta^i)$ be the probability distribution according to which the policymaker selects the initial value of its coefficients.

We intentionally choose a rather naive model selection protocol. One could easily consider a more elaborate and realistic search process. For example, one could place more weight on models with smaller historical mean-squared errors, as was done in our first example. However, it will become apparent that our main conclusions are independent of the details of the selection process. All that we require is a broad, exhaustive search that excludes no model or coefficient vector. While this may engender highly unrealistic model switching dynamics, as poorly fitting parameter vectors are quickly rejected, it does not influence our inferences about the robustness of self-confirming equilibria.

4. ANALYSIS

4.1. Representation as a Stochastic Recursive Algorithm (SRA). It is useful to begin by collecting together the model's equations. We first have a set of model update equations.

$$\hat{\beta}_n^i = \hat{\beta}_{n-1}^i + \eta(R_{n-1}^i)^{-1}\Lambda_n^i \quad (4.14)$$

$$\Lambda_n^i = \Phi_{n-1}^i(x_n - \Phi_{n-1}^{i'}\hat{\beta}_{n-1}^i) \quad (4.15)$$

$$R_n^i = R_{n-1}^i + \eta(\Phi_{n-1}^i\Phi_{n-1}^{i'} - R_{n-1}^i) \quad (4.16)$$

Through feedback, the actual DGP is determined.

$$x_n = A_1(s_{n-1}, \beta_{n-1})x_{n-1} + B_1(s_{n-1}, \beta_{n-1})z_{n-1} + C_1(s_{n-1}, \beta_{n-1})u_n + D_1(s_{n-1}, \beta_{n-1})\epsilon_{1n} \quad (4.17)$$

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

Models are tested by forming recursive LM test statistics

$$\theta_n^i = \theta_{n-1}^i + \eta[\Lambda_n^i \Sigma_{in}^{-1/2} (R_n^i)^{-1} \Sigma_{in}^{-1/2} \Lambda_n^i - \theta_{n-1}^i] \quad (4.18)$$

where Σ_{in} is an $s \times s$ diagonal matrix containing the squared residuals from model i .

In the event the current model is rejected, a new model is selected according to the probability distribution $\pi = (\pi^i)$ over \mathcal{M} , along with a new coefficient vector drawn according to $F^i(\beta^i)$.

These ingredients induce a Markov chain over \mathcal{M} . Let $p_n \in \Delta^{m-1}$ be the time- n probability distribution over models, and let \mathcal{P}_n be an $m \times m$ state transition matrix of

s_n , 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

$$p'_{n+1} = p'_n \mathcal{P}_n \quad (4.19)$$

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

$$\mathbf{P}[\theta_n^i \leq \bar{\theta}] + \mathbf{P}[\theta_n^i > \bar{\theta}] \cdot \pi^i \quad (4.20)$$

and the off-diagonals are given by

$$\mathbf{P}[\theta_n^i > \bar{\theta}] \cdot \pi^j \quad (4.21)$$

where $\bar{\theta}$ is the common test threshold.

Reflecting time-scale separation, it is useful to decompose \mathcal{P}_n into two parts, highlighting the rate at which transitions arise:

$$\mathcal{P}_n = \mathbf{I} + O(\eta, n) \mathbf{Q}_n \quad (4.22)$$

where $\mathbf{Q} = (q_{ij})$ is the generator of a continuous-time Markov chain, and

$$\lim_{\eta \rightarrow 0} O(\eta, n) = 0 \quad \forall n.$$

Note that $\forall i \neq j$, $q_{ij} > 0$ and $q_{ii} < 0$ so that $\sum_j q_{ij} = 0$. $O(\eta, n)$ captures the rate at which a given model is rejected. It plays an important role in the ensuing analysis.

4.2. ODE Approximation of Model Revisions. To approximate the sample path of $\beta_n = (\beta_n^i)$, it is instructive to calculate the speed at which the policymaker switches between models. Following Dupuis and Kushner (1989), the rejection probability in a small neighborhood of a stable self-confirming equilibrium β_e^i is an exponential function of η .

Proposition 4.1. $\forall i, \forall T, \forall \rho > 0, \exists V_i^*(\rho, T) > 0$ such that

$$\mathbf{P}\left(\exists \tau \in (0, T), |\hat{\beta}_\eta^i(\tau) - \beta_e^i| > \rho \mid \hat{\beta}_\eta^i(0) = \beta_e^i\right) \leq e^{-V_i^*(\rho, T)/\eta}.$$

Next, since the test statistic is a continuous function of β^i , we can appeal to the Contraction Principle (Dembo and Zeitouni (1998)) to re-write the above statement in terms of the test statistics.

Corollary 4.2. $\forall i, \forall T, \forall \bar{\theta} > 0, \exists S_i^*(\bar{\theta}, T) > 0$ such that

$$\mathbf{P}\left(\exists \tau \in (0, T), \theta_\eta^i(\tau) > \bar{\theta} \mid \theta_\eta^i(0) = 0\right) \leq e^{-S_i^*(\bar{\theta}, T)/\eta}$$

where θ_η^i is the continuous time process obtained by linear interpolation of $\{\theta_n^i\}$.

Thus, for a sufficiently small $\eta > 0$, the probability that model i is rejected around its self-confirming equilibrium is of order $e^{-S_i^*(\bar{\theta}, T)/\eta}$.

Lemma 4.3.

$$\mathbf{P}\left(\theta_\eta^i(\tau) > \bar{\theta} \mid \hat{\beta}_\eta^i(0) = \beta_e^i\right) = O(e^{-S_i^*(\bar{\theta}, T)/\eta})$$

Although Corollary 4.2 and Lemma 4.3 are proved for the initial condition being set at β_e^i , by exploiting continuity of the logarithmic moment generating function, one can repeat the same logic for any initial condition β_0^i within a small neighborhood of β_e^i .

Lemma 4.4. $\forall i \in \mathcal{M}, \exists \rho_i > 0$ such that $\forall \hat{\beta}_\eta^i(0)$ satisfying $|\hat{\beta}_\eta^i(0) - \beta_e^i| < \rho_i$,

$$\mathbb{P}\left(\theta_\eta^i(\tau) > \bar{\theta} \mid \hat{\beta}_\eta^i(0)\right) = O(e^{-S_i^*(\bar{\theta}, T)/\eta})$$

Let $N_{\rho_i}(\beta_e^i)$ be a ρ_i neighborhood of β_e^i . For the moment, let us examine the evolution of the stochastic recursive system described by (4.14)-(4.21), conditioned on the event that $s_n = i$ and $\hat{\beta}_n^i \in N_{\rho_i}(\beta_e^i)$.

By Lemma 4.4, we know that a transition to another model is a very unlikely event, whose probability vanishes at an exponential rate. Exploiting this feature, we can invoke an approximation theorem of Yin and Krishnamurthy (2005) to approximate the sample path starting at $\hat{\beta}_n^i \in N_{\rho_i}(\beta_e^i)$ by the ODE (3.12).

Proposition 4.5. Fix $\hat{\beta}_0^i \in N_{\rho_i}(\beta_e^i)$, and $s_0 = i$. Then,

$$\lim_{\eta \rightarrow 0} \hat{\beta}_\eta^i(\tau) \Rightarrow \beta^i(0) + \int_0^\tau h^i(\beta^i) ds$$

where \Rightarrow denotes weak convergence.

Since in a small neighborhood of β_e^i the rejection probability of model i vanishes at an exponential rate, the log of the mean duration that β_n^i remains in a small neighborhood of β_e^i is proportional the rate function $S_i^*(\rho_i, T)$ (Dembo and Zeitouni (1998)).

Lemma 4.6. Fix $\hat{\beta}_0^i \in N_{\rho_i}(\beta_e^i)$, and $s_0 = i$. Let τ^η be the first time that model i is rejected. Then,

$$\lim_{\eta \rightarrow 0} \eta \log \mathbb{E}[\tau^\eta \mid \hat{\beta}_0^i \in N_{\rho_i}(\beta_e^i)] = S_i^*(\rho_i, T)$$

and $\forall \delta > 0$,

$$\lim_{\eta \rightarrow 0} \mathbb{P}\left(e^{\frac{S_i^*(\rho_i, T) + \delta}{\eta}} > \tau^\eta > e^{\frac{S_i^*(\rho_i, T) - \delta}{\eta}}\right) = 1.$$

If $\hat{\beta}_n^i$ visits a small neighborhood of β_e^i while $s_n = i$, then the model survives testing for an exponentially long period of time, while staying within a small neighborhood of the self-confirming equilibrium.

We have yet to characterize the dynamics of (4.14)-(4.21) outside small neighborhoods of each model's self-confirming equilibrium. The following result says that eventually, due to our assumptions of exhaustive search and global stability of SCE, the policymaker's model will come arbitrarily close to a SCE.

Proposition 4.7. $\forall \delta \in (0, 1), \exists T^*$ and $\eta^* > 0$ such that

$$\mathbb{P}\left(\exists n \leq \lceil T^*/\eta \rceil : s_n = i, \hat{\beta}_n^i \in N_{\rho_i}(\beta_e^i)\right) \geq 1 - \delta.$$

Proof. Recall that the set of feasible coefficients is compact and convex, and that in response to a model rejection, the policymaker selects every model with positive probability. Since the set of coefficients around a stable self-confirming equilibrium has positive measure, the decision maker must visit this neighborhood within a finite number of periods. Specifically, $\forall \delta_1 > 0, \exists K_1$ such that

$$\mathbb{P}\left(\exists i, \exists n \leq K_1 : \hat{\beta}_n^i \in N_{\rho_i}(\beta_e^i)\right) \geq 1 - \delta_1. \quad (4.23)$$

Of course, if models are rejected quickly, the policymaker switches models frequently. Although it is possible that his model survives the test even when $\hat{\beta}_n^i \notin N_{\rho_i}(\beta_e^i)$, since the SCE is globally stable and the policymaker continuously updates $\hat{\beta}_n^i$ while it survives the test, the coefficient vector β_n^i must converge to it. Also, since the domain of the coefficients is compact, it will take a finite number of iterations to visit a small neighborhood of β_e^i .

To formalize this observation, let $\tau_i(\hat{\beta}_0^i)$ be the time it takes for $\hat{\beta}_\eta^i(\tau)$ to evolve from $\hat{\beta}_0^i$ to $N_{\rho_i/2}(\beta_e^i)$, if it evolves according to the ODE (3.12). Since the trajectory is a continuous function of the initial condition, \mathcal{B}^i is compact, and the SCE is globally stable,

$$\bar{\tau}_i = \sup_{\hat{\beta}_0^i \in \mathcal{B}^i} \tau_i(\hat{\beta}_0^i) < \infty$$

Let

$$\bar{\tau} = \max_i \bar{\tau}_i.$$

Since the sample paths of the stochastic recursive algorithm converge to the trajectory of (3.12), $\forall \delta_2 > 0$, $\exists \eta_2 > 0$ such that $\forall \eta \in (0, \eta_2)$, $\forall i$, $\forall \hat{\beta}_0^i \in \mathcal{B}^i$

$$\mathbb{P} \left(\exists \tau \leq \bar{\tau} : \hat{\beta}_\eta^i(\tau) \in N_{\rho_i/2}(\beta_e^i) \right) \geq 1 - \delta_2 \quad (4.24)$$

Choose $\delta_1 = \delta_2$, $\delta = \delta_1 + \delta_2$, and $T^* = K_1 \bar{\tau}$. The conclusion then follows from combining (4.23) and (4.24). \square

From Lemma 4.6 and Proposition 4.7, we conclude that as $\eta \rightarrow 0$, the (fictitious) time the system takes to reach a small neighborhood of some model's self-confirming equilibrium is uniformly bounded by $\bar{\tau}$, while the duration around it increases exponentially as $\eta \rightarrow 0$. Moreover, this duration is also exponentially proportional to the rate function S_i^* associated with model i 's self-confirming equilibrium.

Given an integer N , define model i 's empirical occupancy measure

$$\pi(i, N) = \frac{1}{N} \#\{s_n = i\}$$

as the proportion of periods during which model i is used. Define i^* as the model which has the largest rate function

$$S_{i^*}^*(T, \beta_e^{i^*}) \geq S_j^*(T, \beta_e^j) \quad \forall j.$$

We call $\beta_e^{i^*}$ a robust self-confirming equilibrium. Our main result is the following, which follows in an obvious way from the previous results.

Theorem 4.8.

$$\lim_{N \rightarrow \infty} \pi(i^*, N) = 1$$

On the one hand, this result is quite general, as it allows us to identify a dominant model independently of the details of model selection and experimentation. On the other hand, it is quite limited, since without these details one cannot make useful predictions about the real-time dynamics of model selection.

4.3. Discussion. We have defined a robust self-confirming equilibrium as an outcome generated by a model which survives specification testing longer than any other model in a given model class. 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 entropy, and then, from either Stein's lemma (classical) or Chernoff bounds (Bayesian), we know that relative entropy is connected to detection error probabilities. In particular, larger relative entropy should make it easier to discriminate between a model and the true DGP. That is, larger relative entropy reduces the probabilities of Type I and Type II errors. Why then are models with large rate functions *more* robust against specification tests?

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).

It is natural to ask what the relationship is between the dominant model and the actual DGP. For example, if the true model is in \mathcal{M} , will it be dominant? Interestingly, with feedback this may not be the case. Feedback could accelerate the escape of a true model. However, without feedback, the policymaker's role reverts to that of an outside econometrician, and we get the usual results (White (1982)).

Proposition 4.9. *Suppose the DGP is exogenous. If the true model is in \mathcal{M} then it is dominant. If the true model is not in \mathcal{M} , the dominant model will be the model that is closest in terms of KLIC distance.*

Since our primary interest is in the endogenous data case, we shall not formally prove this. However, it follows easily from two facts. First, with random model selection and a linear model class there is no essential difference between model uncertainty and parameter uncertainty, and our assumptions imply the parameter space is strictly convex. Second, since the environment is Gaussian, recursive least squares becomes equivalent to maximum likelihood.

5. ANOTHER EXAMPLE

At this point it is useful to consider another example, designed to illustrate some additional features of model validation that were not present in the cobweb example of section 2. In this example, the agent uses his model to solve a control problem, which creates an additional source of feedback. Moreover, the misspecification here is more subtle than simply omitting relevant explanatory variables. We shall see that by validating his model, the agent learns to do the right thing, thus vindicating Sims (1982). At the same time, however, we shall also see that the ability to conduct specification searches opens the door to more insidious outcomes, in which mistaken beliefs can be ex post rationalized by the validation process.

5.1. 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,




FIGURE 3. Simulated Inflation Path in Sargent’s Conquest Model

and then occasional sharp inflation stabilization. As noted by Cho, Williams, and Sargent (2002) (CWS), this cycle represents the interplay between the model’s mean dynamics and its 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 trade-off, 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.

5.1.1. 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 3 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 4.




FIGURE 4. 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.¹⁵ 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 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.

5.1.2. *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

¹⁵Again, leave aside for the moment the naivete associated with repeatedly applying a t -test with a fixed threshold.

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 5




FIGURE 5. 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’.

5.1.3. *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).¹⁶ So let’s now suppose the Bank 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 6 contains the results¹⁷

¹⁶Presumably, the researcher has access to a time machine, as these methods were not developed until recently.

¹⁷A 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 .




FIGURE 6. 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 well into the 1980s, although of course during the 1980s things were moving in reverse, with oil prices and unemployment both falling.

What is interesting about Figure 6 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 6, 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 7 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 7. 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.

5.2. Learning To Do The Right Thing. We just 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 5.1.2 never gets hired, but instead the Bank follows the model validation procedure we outlined in Sections 3 and 4. 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 i.i.d 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. 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.

6. 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 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

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.

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).

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. Second, 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 selection. Policymakers do not randomly select models. It would be interesting to study the effects of more realistic model selection strategies. Although this would not affect the main conclusions of this paper, it would open the door to a real-time analysis of model selection dynamics, with potentially interesting empirical implications. Third, it would be useful to provide a more complete characterization of the dominant model by studying in more detail the links between large deviation rate functions and relative entropy. 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 (2009).

REFERENCES

- AN, S., AND F. SCHORFHEIDE (2007): "Bayesian Analysis of DSGE Models," *Econometric Reviews*, 26, 113–72.
- 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., 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.
- BURNHAM, K., AND D. ANDERSON (2002): *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*. Springer, 2nd edn.
- CHO, I.-K., AND K. KASA (2009): "Recursive Model Selection with Endogenous Data," 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., AND T. J. SARGENT (2005): "The Conquest of U.S. Inflation: Learning and Robustness to Model Uncertainty," *Review of Economic Dynamics*, 8, 528–63.
- 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.
- DUPUIS, P., AND H. J. KUSHNER (1989): "Stochastic Approximation and Large Deviations: Upper Bounds and w.p.1 Convergence," *SIAM Journal of Control and Optimization*, 27, 1108–1135.
- ENGLE, R. F. (1984): "Wald, Likelihood Ratio, and Lagrange Multiplier Tests in Econometrics," in *Handbook of Econometrics, Volume II*. Elsevier Science Publishers.
- 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 (2007a): "Recursive Robust Estimation and Control Without Commitment," *Journal of Economic Theory*, 136, 1–27.
- (2007b): *Robustness*. Princeton University Press.
- KANDORI, M., G. MAILATH, AND R. ROB (1993): "Learning, Mutation and Long Run Equilibria in Games," *Econometrica*, 61, 27–56.

- 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, R. E., AND T. J. SARGENT (1981): *Rational Expectations and Econometric Practice*. University of Minnesota Press.
- MARCET, A., AND T. J. SARGENT (1989): "Convergence of Least Squares Learning Mechanisms in Self Referential Linear Stochastic Models," *Journal of Economic Theory*, 48, 337–368.
- 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.
- 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.
- YIN, G. G., AND V. KRISHNAMURTHY (2005): "LMS Algorithms for Tracking Slow Markov Chains With Applications to Hidden Markov Estimation and Adaptive Multiuser Detection," *IEEE Transactions on Information Theory*, 51(7), 2475–91.