

LEARNING AND MODEL VALIDATION

IN-KOO CHO AND KENNETH KASA

ABSTRACT. This paper 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 better-fitting model. If his model passes the test, he uses it to formulate and 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 at all well understood. We make progress on this problem by relating it to a problem that *is* well understood. In particular, we relate it to the dynamics of constant-gain stochastic approximation algorithms. Doing this enables us to appeal to well known results from the large deviations literature to help us understand the dynamics of testing and model revision. We show that as the agent applies an increasingly stringent specification test, the large deviation properties of the discrete model validation dynamics converge to those of the continuous learning dynamics. This sheds new light on the recent constant-gain learning literature.

JEL Classification Numbers: C120, E590

1. INTRODUCTION

Since the days of Frisch, Tinbergen, and Haavelmo, econometrics has held out the hope of improving government policy, especially macroeconomic policy. Initially, efforts focused on issues of simultaneity and identification, since these problems had not yet been confronted by the experimental data of the natural sciences. By the early 1970s, these problems were largely resolved, and applied econometricians could ply their trade using a sophisticated toolbox of instrumental variables and two- and three-stage least squares estimators.

Unfortunately, just as these methods began to seem routine, Lucas [30] and Sims [39] advanced alternative, but equally devastating, critiques of the Cowles Commission methodology. Lucas argued that basic dynamic economic theory predicted that econometric models would not be invariant to changes in government policy. Policy changes would change agents' decision rules, and this would change the model. In this case, using a model to formulate a *new* policy is invalid, unless it properly accounts for the reactions of the private

Date: September 17, 2005.

We would like to thank Lars Hansen and Tom Sargent for helpful discussions. Financial support from the National Science Foundation (SES-0004315, ECS-0523620) 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.

sector. Cowles Commission methods did not do this.¹ Sims, on the other hand, did not worry about regime changes. He argued that the notion of a regime change was fraught with conceptual difficulties, associated with observing unexpected events, and even if such things could be well defined, they did not happen very often in practice, and hence, were of limited importance to the task of implementing normal science. Instead, Sims focused on the dubious nature of the Cowles Commission identification strategy, based on large numbers of zero restrictions.

In principle at least, both of these critiques have been overcome. The Lucas critique suggested an alternative identification strategy, based on cross-equation restrictions (see Lucas and Sargent [29]). The Sims Critique was largely overcome by Sims himself, in his development of VAR methods. Despite these successes, it is probably fair to say that a cloud of skepticism still hangs over most macroeconomic research.

Why is this? One possible answer stems from a crucial assumption which, until quite recently, has been maintained by virtually all macroeconomic research. Responding to the Lucas Critique requires assumptions about the models that agents use when formulating their plans. The maintained assumption of current macroeconomic research is that the government policymaker “knows the model”. Not only that, he knows the private sector knows the model, and they know he knows the model, and so on, *ad infinitum*. That is, both the government and the private sector have a correctly specified model, and this model is common knowledge.²

Economists have always been uncomfortable with this assumption.³ However, it has only been recently that much progress has been made in weakening it. Our approach is inspired by Sargent’s [36, 37] work on *bounded rationality*. This approach combines the following elements: (1) There is a decisionmaker that takes actions based on a possibly misspecified model, (2) These actions influence the true data-generating process. This is descriptive of many *large agent* situations, like government policymaking. It also wreaks havoc when attempting to apply classical statistical methods, which are based on the assumption of an exogenous data-generating process. Instead, we are confronted with a self-referential dynamic system [31] in which the data generating process interacts with the decision maker’s belief formation process. As a result, all the classical convergence and consistency theorems go out the window, and we must devise new methods that incorporate the following features: (3) The decision maker is aware of potential model misspecification, and tries to detect it, in real-time, using best practice econometric methods; (4) If the current model is rejected, he formulates a new model that is more in line with the data, and finally; (5) If the current model is not rejected, he formulates and implements a policy based on the provisional assumption that the current model is correctly specified and will not change in the future. This last ingredient puts us squarely in the bounded rationality camp, since the decisionmaker fails to recognize and respond to his own influence over future data. He learns, but purely in a passive, retrospective way. Of course, it also

¹However, Jacob Marschak [32], a Cowles founding father, was aware of the problem.

²During the past couple of years there has been a flood of research on model uncertainty and robust policies. We relate our approach to this research as we go along.

³See Frydman and Phelps [20] for an early collection of articles expressing doubts about the common knowledge assumptions of the Rational Expectations Hypothesis

side-steps the conceptual difficulties of a Bayesian approach (see, e.g., Bray and Kreps [4]).

Confronted individually, each of these features can be handled. For example, due to the work of Vuong [41] and Hansen and Sargent [21], econometricians now know how to test and compare misspecified models. Due to the work of Chu, Stinchcombe, and White [11] econometricians know how to detect model changes in real-time. Due to the work of Marcet and Sargent [31] economists know how to analyze self-referential dynamic systems. Unfortunately, when you combine all these elements, there are few, if any, results.

Our paper attempts to make progress on this difficult problem. We do this by relating it to a *different* problem with known features. Specifically, we show that in a certain limiting sense the dynamics produced by a testing-and-model-validation process are equivalent to the dynamics produced by a continuously revised, constant-gain stochastic approximation algorithm. Given this, we can then approximate our model validation dynamics by appealing to the recent results of Williams [42] and Cho, Williams, and Sargent [10]. These papers apply large deviation methods to characterize the escape dynamics of adaptive learning models. Escape dynamics are proving to be a useful way to model a wide-range of markov-switching data, such as inflation stabilizations (Sargent [37]) and recurrent currency crises (Kasa [23] and Cho and Kasa [8]).

The convergence between these two processes occurs as the decisionmaker applies an increasingly stringent specification test. When testing, we assume he employs a generalized relative entropy test (see Dembo and Zeitouni [17], ch. 7), which is known to be optimal in the Neyman-Pearson sense, even for quite general alternatives. For this reason, Zeitouni and Gutman [43] call it a ‘universal hypothesis test’. The use of a relative entropy test is also quite natural here, given our focus on large deviations, since we know from Stein’s lemma that there is a close connection between Type I and Type II error rates and relative entropy, and we know from Sanov’s Theorem that there is a close connection between relative entropy and large deviation rate functions.⁴ The test accepts the current model if and only if the relative entropy between the data and the model does not exceed a given threshold, say $\rho > 0$. We do not have to say much about this threshold, since we focus on the case $\rho \downarrow 0$. More generally, one could specify ρ to optimally trade-off Type I and Type II errors.⁵

When his model is rejected we assume the decisionmaker builds a new model using maximum likelihood estimation. When ρ is large, these model revisions occur infrequently, and as a result, the reference model path tends to be rather ‘jerky’. However, we show that as $\rho \downarrow 0$ the discrete model revision dynamics converge in a very strong sense to the continuous recursive learning dynamics. In particular, both the testing dynamics and the learning dynamics induce probability distributions over sample paths. Our convergence result implies that these two probability distributions coincide not only in the ‘center’ of the distribution, *but also in the tails*. From this, we can conclude that the two processes share the same large deviation properties (e.g., escape routes and escape times) as $\rho \downarrow 0$.

⁴See Cover and Thomas [16] for a nice discussion of Stein’s lemma. Dembo and Zeitouni [17] provide a thorough discussion of Sanov’s Theorem.

⁵See, e.g., Lai and Shan [28]. They point out that under certain conditions relative entropy tests are intimately related to more familiar CUSUM tests.

This asymptotic equivalence result sheds new light on the rapidly growing constant-gain learning literature. The early least-squares learning literature focused on the question of asymptotic convergence to a Rational Expectations equilibrium.⁶ This is an inherently nonstationary problem, and is mainly of theoretical interest. To explain observed time series, there are great advantages to studying stationary equilibria. Constant gain learning algorithms produce stationary equilibria. The recent constant gain learning literature has shown that when agents suspect that the model they are learning about may be changing over time, and they respond to their suspicions by placing more weight on recent data, then persistent cyclical dynamics can sometimes result. Interestingly, while these dynamics confirm the agents' suspicions, they do so in an entirely self-referential way, since they only occur *because of* the agents' suspicions.

One common way to motivate constant gain learning is to highlight its connections to more familiar Kalman filtering approaches. As noted by Sargent [37], and discussed in more detail by Sargent and Williams [38], constant gain recursive learning algorithms are (nearly) equivalent to the Kalman filter when priors are based on 'slow' random walk parameter drift. Drawing this connection is useful, because it reveals a potential problem with constant gain learning models. In particular, the small-noise/random-walk specification of parameter drift in a sense 'commits' the agent to a constant rate of parameter drift, when in fact, during escape episodes, the parameters may be changing quite rapidly. This raises the question of whether we are inappropriately tying the hands of the agents in our models.⁷

A model validation approach potentially avoids this inconsistency. Rather than adopt the somewhat artificial assumption that agents update their models each period at the same rate, we can instead adopt the more realistic assumption that agents *monitor* their models continuously, but only abandon them if they go sufficiently far off track. A time invariant rejection threshold will automatically generate a greater rate of model revision during times of economic turbulence. However, as we discuss later, a model validation approach does face a similar issue, since we must decide how to use past data when constructing our test statistics. The statistically optimal rate of discount may well be time varying. We briefly touch upon this issue at the end of the paper, in the context of robust hypothesis testing [43, 33, 34].

Another payoff from our result is that it helps us to understand a puzzling feature of macroeconomic policy. Recent work by Cogley and Sargent [13] suggests that the Fed actually began to suspect the short-term Phillips curve was misspecified by the mid-1970s, several years *before* the Volcker disinflation. Cogley and Sargent attribute this policy inertia to Bayesian model uncertainty. The Fed stuck to a high inflation policy because it risk-dominated the policy implied by the better-fitting model. Our paper provides an alternative interpretation of policy inertia - even if a model is rejected, it is typically the case that a similar model is the best fitting model. That is, even though a tail-sensitive test, like our relative entropy test, may detect a change, estimators like least squares or

⁶The treatise by Evans and Honkapohja [19] provides a definitive survey of this literature.

⁷Note, this issue is related to a recent debate between Sims [40] and Cogley and Sargent [12]. Sims argues that heteroskedasticity, which is not properly accounted for in agents' models, can fool them into thinking their models are changing. Recent empirical work by Cogley and Sargent [15] suggests a role for both.

maximum likelihood, which focus on fitting the center of the distribution, will typically dictate a modest model revision. Drastic policy changes only take place at self-confirming equilibria, where model rejections are surprising, and therefore, informative.

The remainder of the paper is organized as follows. Section 2 motivates the analysis by outlining a version of the canonical Phillips curve example of Sargent [37]. Section 3 contains a detailed description of both the recursive learning dynamics and the model validation dynamics. Various assumptions and regularity conditions are also discussed. Section 4 proves that the H-functionals of the two processes converge to each other as $\rho \downarrow 0$. From Kushner [27], this is sufficient to establish convergence of the escape dynamics. Section 5 discusses extensions, and section 6 concludes with some suggestions for future research.

2. AN EXAMPLE

Although we formulate our results in a general setting, it is useful to begin with a concrete example. A general treatment of the problem involves a lot of notation and terminology, and it will be helpful to have observable counterparts in mind as we go along. To do this we borrow heavily from Sargent [37] and Cho, Williams, and Sargent [10] (henceforth denoted CWS).

2.1. Recursive Learning. Consider the following example of Sargent [37], in which unemployment u_t and inflation y_t are assumed to evolve according to the expectation-augmented Phillips curve:

$$(2.1) \quad u_t = u^* - \theta(y_t - x_t^e) + v_{1t},$$

where $\theta > 0$, x_t^e is the private sector's forecast of inflation, u^* is the natural rate of unemployment, and v_{1t} is white noise. The realization of actual inflation y_t is determined by the government's inflation target x_t and a nominal shock v_{2t} :

$$(2.2) \quad y_t = x_t + v_{2t}$$

where v_{2t} is also white noise. Sargent [37] and CWS [10] impute a subtle form of misspecification to the government, which misinterprets the role of private sector beliefs. In particular, in place of the expectations-augmented Phillips curve, the government mistakenly believes in a Keynesian short term Phillips curve⁸.

$$(2.3) \quad u_t = \gamma_0 + \gamma_1 y_t$$

This misspecification injects parameter drift into the government's approximating model, arising from the evolving beliefs of the private sector. The government responds to this drift by using a discounted recursive least squares algorithm to update the coefficient of its model. Letting $\gamma = (\gamma_0, \gamma_1)^T$ we have

$$(2.4) \quad \hat{\gamma}_{t+1} = \hat{\gamma}_t + aR_t^{-1}\phi_t e_t, \quad R_{t+1} = R_t + a(\phi_t \phi_t^T - R_t), \quad t \geq 0,$$

where the regression vector is $\phi_t = (1, y_t)^T$, the innovation is $e_t = u_t - \hat{\gamma}_{0t} - \hat{\gamma}_{1t}y_t$, and a is the gain sequence. Some form of discounting is critical since one hopes that the estimates

⁸This specification error is subtle in a sense made precise by Sargent [35]: given any historical data on y_t and u_t , the two models (2.1) and (2.3) are *observationally equivalent*

will approximate the current best model in the parameterized model class in a changing environment. The target inflation rate is selected by using the perceived model in (2.3) to solve the so-called Phelps problem,

$$\min_{x_t} \mathbf{E}[u_t^2 + y_t^2]$$

The solution gives the target inflation rate as a function of the government's model, parametrized by $\hat{\gamma}_t$. In this example, we invoke Sargent's "Fed Watcher" assumption, meaning that the private sector knows the government's inflation target:

$$x_t^e = x_t.$$

This allows us to focus on the behavior of the government, and can be replaced by a much milder assumption. Note that the distribution of y_t is affected by three elements: (1) the government's model $\hat{\gamma}_t$, (2) the state observed by the government, and (3) any relevant states the government overlooks (in this case, the private sector's expectation). Let X_t be the observed state and Z_t be any state which the government is not aware of, or simply, fails to include in (2.3).⁹ We can write

$$(2.5) \quad y_t = b^r(\hat{\gamma}_t, X_t, Z_t).$$

Following Sargent [37], we can calculate the Kydland and Prescott outcome, which corresponds to the Nash equilibrium or the self-confirming equilibrium where

$$(2.6) \quad \gamma^e = (\gamma_0^e, \gamma_1^e) = (u^*(1 + \theta^2), -\theta)$$

and the government sets the inflation target as

$$x_t = \theta u^*.$$

Another important outcome is the Ramsey outcome which corresponds to the Friedman rule where

$$\gamma^r = (\gamma_0^r, \gamma_1^r) = (u^*, 0)$$

and the government is committed to

$$x = 0.$$

One can show that as $a \rightarrow 0$, the sample path induced by the recursive learning dynamics can be approximated by a trajectory induced by an ordinary differential equation (ODE)

$$\dot{\gamma} = \varphi(\gamma)$$

which has (2.6) as a unique stable stationary point. By "approximated", we mean the sample paths induced by the two dynamics are close in the topology of weak convergence as the gain $a \rightarrow 0$ in (2.4).

A simulation of a typical sample path of $\{y_t\}$ based on these specifications is shown in Figure 1. Although this is a highly stylized model, the overall dynamics capture the nonlinear aspect of post WWII U.S. inflation, i.e., persistent increases followed by sharp decreases. The simulation also suggests that it might be dangerous to conclude that U.S.

⁹We interpret the missing state variables Z_t broadly. Suppose the government becomes aware of the fact that realized inflation *is* affected by expected inflation, but does not correctly specify the functional form for the private sector's expectation formation process. In this case, Z_t would capture the gap between the true expectations of the private sector and the expectations calculated by the government according to its misspecified model.

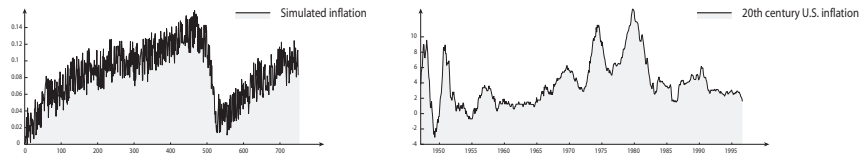


FIGURE 1. The right panel is the moving average of monthly C.P.I. inflation (all items), and the left panel is simulated inflation based on the recursive learning model.

inflation has been conquered, despite the fact that inflation has been under control for the past 25 years [37].¹⁰ That is, Figure 1 reveals the need to consider large deviation properties when evaluating the long run performance of policies [37]. With occasional excursions away from the self-confirming equilibrium, the overall average inflation rate is significantly lower than the self-confirming equilibrium rate.

While recursive learning models offer an elegant explanation of these complex dynamics, this approach has been subject to criticism since the original work of Bray [3]. Learning models are motivated by misgivings about the model specification assumptions of the Rational Expectations Hypothesis: even if agents do not know the exact model, they at least know it up to some finite dimensional parametric class, and moreover, they all have the *same* model. Learning models allow agents to entertain *diverse* and *approximate* models, and to update their models in light of new evidence. Still, it is rather discomfoting to assume that agents use the same misspecified model over time, without even a drop of suspicion about the parametric validity of the model. Once the modeler endows him with a parametric model, the agent is only allowed to update the parameters to fit the data.¹¹ As noted in the Introduction, this is particularly disturbing in this constant gain example, because the data on inflation indicates that the underlying state may not be stationary. In this case, smart agents may decide to look for a new parametric model.

A natural question is whether key qualitative and quantitative features of the learning dynamics continue to hold if the model is subject to specification tests prior to its use for decision making. We now compare two distinct settings: (1) The conventional least squares learning model just discussed, and (2) An alternative model in which the decision maker continually tests the “validity” of his model. One might expect the resulting *validation dynamics* to be very different from the dynamics found in the conventional model. In fact, the reverse is true, which provides a powerful response to the above criticism of recursive learning models.

2.2. Model validation. The key step to formalizing the idea of model validation is to formulate a method that gives the government some sense of the reliability of the current model based on the current estimate $\hat{\gamma}_t$. In this example, we impose two restrictions in order to illuminate the key ideas.

¹⁰Further details along with a formal analysis can be found in [10].

¹¹It should be noted that there are small literatures devoted to relaxing this assumption. See, e.g., Chen and White’s [7] work on nonparametric recursive learning, and Brock and Hommes [6] work on forecast model selection.

- A1. $v_{i,t} \sim N(0, \sigma_i^2)$, and the policy maker knows the correct distribution.
A2. the policy maker's model is confined to the parametric class (2.3).

These assumptions greatly simplify the computation of the test statistics. Because the policy maker knows the distribution, we side-step a potentially important source of model uncertainty.¹² We shall return to distributional uncertainty after we present the main results. The second assumption delineates the class of models of the government, which is the class of linear regression models with 2 fixed explanatory variables. In principle, the number of explanatory variables could be endogenously determined through a model selection process.

Let \mathcal{M}_t be the test statistics computed from the data. While one can consider different kinds of specification tests, we opt for the relative entropy criterion for two reasons. First, for a broad class of models, a threshold relative entropy test is optimal in the sense of Neyman-Pearson [43, 17]. In fact, relative entropy and likelihood ratio tests are widely used as specification tests (e.g., [41]). Second, relative entropy is intimately linked to the Kullback-Leibler information criterion for model selection. This enables us to formulate the testing and selection of models in a unified framework.

Since we restrict our focus to linear Gaussian models, \mathcal{M}_t can be completely represented by the sample mean $\hat{\mu}_t$ and the sample covariance matrix $\hat{\Sigma}_t$, which can be estimated recursively as

$$\begin{aligned}\hat{\mu}_{t+1} &= \hat{\mu}_t + a_g \left(\begin{bmatrix} u_t \\ y_t \end{bmatrix} - \hat{\mu}_t \right) \\ \hat{\Sigma}_{t+1} &= \hat{\Sigma}_t + a_t \left(\begin{bmatrix} u_t \\ y_t \end{bmatrix} - \hat{\mu}_t \right) \left[[u_t \quad y_t] - \hat{\mu}_t^T \right] - \hat{\Sigma}_t\end{aligned}$$

where $[\cdot]^T$ is the transpose of $[\cdot]$.

Suppose that the government's model is $\tilde{\gamma}_k = (\tilde{\gamma}_{0,k}, \tilde{\gamma}_{1,k})$. Let $\mathcal{M}_{\tilde{\gamma}_k}$ be the probability distribution over (u_t, y_t) when the government takes an action based on $\tilde{\gamma}_k$. Define the relative entropy between the model and the data as

$$I(\mathcal{M}_t \| \mathcal{M}_{\tilde{\gamma}_k}) = \int \frac{d\mathcal{M}_t}{d\mathcal{M}_{\tilde{\gamma}_k}} d\mathcal{M}_t.$$

Let $\mu_{\tilde{\gamma}_k}$ and $\Sigma_{\tilde{\gamma}_k}$ be the mean and the covariance matrix of $\mathcal{M}_{\tilde{\gamma}_k}$. If both models have Gaussian distribution, we then have

$$I(\mathcal{M}_t \| \mathcal{M}_{\tilde{\gamma}_k}) = \frac{1}{2} \left(-\log \frac{|\hat{\Sigma}_t|}{|\Sigma_{\tilde{\gamma}_k}|} + \text{trace} \left(\Sigma_{\tilde{\gamma}_k}^{-1} \hat{\Sigma}_t \right) - 2 + (\hat{\mu}_t - \mu_{\tilde{\gamma}_k})^T \Sigma_{\tilde{\gamma}_k}^{-1} (\hat{\mu}_t - \mu_{\tilde{\gamma}_k}) \right)$$

At the beginning of period t , the government has model $\tilde{\gamma}_k$, characterized by an estimated short term Phillips curve

$$u_t = \tilde{\gamma}_{0,k} + \tilde{\gamma}_{1,k} y_t$$

It then puts $\tilde{\gamma}_k$ to the test:

$$(2.7) \quad H_0: (u_t, y_t) \text{ is generated by } \mathcal{M}_{\tilde{\gamma}_k}.$$

¹²This issue is addressed in [33, 34].

The government *validates* $\tilde{\gamma}_k$ if and only if

$$I(\mathcal{M}_t || \mathcal{M}_{\tilde{\gamma}_k}) \leq \rho$$

for a fixed threshold $\rho > 0$. If H_0 is accepted, the government uses the short term Phillips curve to set the inflation target in period t . If $\tilde{\gamma}_k$ is rejected, then the government builds a new reference model $\tilde{\gamma}_{k+1}$ by solving

$$\min_{\tilde{\gamma}} I(\mathcal{M}_t || \mathcal{M}_{\tilde{\gamma}}).$$

Within the confines of linear Gaussian models, $\tilde{\gamma}_{k+1}$ is precisely the maximum likelihood estimator, which can be calculated using recursive least squares. Then, the government chooses the inflation target in period t based on $\tilde{\gamma}_{k+1}$. We refer to the resulting dynamics of $\tilde{\gamma}_k$ and y_t as the *validation dynamics*.

2.3. Simulations and Observations. Figure 2 reports typical sample paths from two simulations when the threshold for the specification test is $\rho = 0.0015$, which in this case is quite stringent. Not surprisingly, when the threshold is relatively large, the validation dynamics can differ markedly from the learning dynamics, since the parameters under the validation dynamics are evolving much more slowly than under the recursive learning dynamics. However, as Figure 2 reveals, as the government runs a more stringent specification test, the two sample paths become virtually identical.

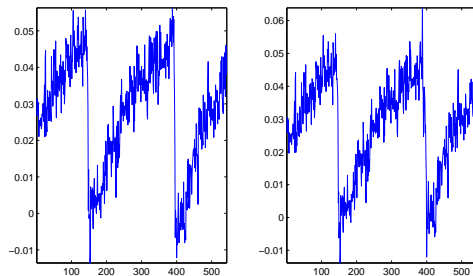


FIGURE 2. Typical sample paths of inflation with more stringent specification tests. The left panel reports the sample path from the validation dynamics, and the right panel depicts the sample path generated by the recursive learning dynamics

We shall formally demonstrate that as the government runs an increasingly tight validation test by reducing the threshold $\rho > 0$, the large deviations properties of the validation dynamics converge to those of the learning dynamics. This asymptotic equivalence provides a new behavioral foundation for recursive learning dynamics.

Validation tests, especially those based on relative entropy, can detect discrepancies between the reference model and the data even in the tails of the distribution. However, maximum likelihood or least squares estimators choose models that fit the data near the center of the empirical distribution. Even if the short term Phillips curve reveals discrepancies from the data in the tails of the distribution, the best fitting model remains

the short term Phillips curve until the state of the economy reaches the self-confirming equilibrium. U.S. data suggests this occurred in early 1980's. The resulting policies exhibit inertia along the convergent path. Even when the short term Phillips curve is rejected by the specification test, the government may appear to cling to the same model with only slightly different parameters, pursuing the same high inflation policy. However, once the economy reaches the self-confirming equilibrium, rejections of the reference model become unlikely events. Thus, conditioned on the rejection, the new reference model can differ significantly from the status quo model, which explains the drastic policy changes that take place around the self-confirming equilibrium.

3. A GENERAL FRAMEWORK

3.1. Recursive Learning. Let $(X, Z) \in \mathbf{X} \times \mathbf{Z} \subset \mathbb{R}^m \times \mathbb{R}^{n-m}$ be a state vector, where X is a vector of observable variables, and Z represents variables which the decision maker does not observe, or simply does not include in his model for any reason. To simplify exposition, assume that \mathbf{X} is compact.¹³ Let \mathbb{P} be the set of all possible models available to the decision maker. We assume that \mathbb{P} is the set of all linear models defined over the elements of X , including lags of X up to $\bar{\ell}$ periods. In this paper, we assume that $\bar{\ell}$ is exogenously given so that we can parameterize a model by a vector β of coefficients in $\mathbb{R}^{m(\bar{\ell}+1)}$. This permits us to represent a model as a vector autoregression:

$$(3.8) \quad \sum_{\ell=0}^{\bar{\ell}} L^\ell \beta_t X_t = \xi_t$$

where L is the lag operator

$$L^\ell \beta_t X_t = \beta_{t-\ell} X_{t-\ell}$$

and ξ_t is the regression residual. Let β be the profile of all regression coefficients.

Given β_t , the decision maker takes an action $b^r(\beta_t, X_t) \in X$, which results in an outcome profile $(y_t, z_t) \in \mathbf{X} \times \mathbf{Z}$ in period t from which only y_t is observed by the decision maker. We assume that $b^r(\cdot)$ obtains by solving an optimization problem. Based on y_t , the decision maker modifies the regression coefficient β_t according to a recursive formula:

$$(3.9) \quad \beta_{t+1} = \beta_t + a\Psi(\beta_t, X_t, Z_t)$$

for $a > 0$, which is the gain sequence. Note that the right hand side includes Z_t , which the decision maker is not aware of. Although Z_t is not observed, it affects the evolution of β_t because y_t is influenced by Z_t . Let

$$(3.10) \quad \dot{\beta} = \bar{\Psi}(\beta)$$

be the associated ordinary differential equation for (3.9):

$$\bar{\Psi}(\beta) = \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T \Psi(\beta_t, X_t, Z_t) \mid \beta_1 = \beta, X_1 = x, Z_1 = z \right].$$

We make a number of assumptions to ensure that (3.10) is a reasonable approximation of the sample paths induced by (3.9) as $a \rightarrow 0$.

¹³The case where \mathbf{X} is unbounded can be handled if the probability distribution of X has sufficiently thin tails.

- Assumption 3.1.** (1) If $\exists K \subset \mathbb{R}^m$ which is convex and compact so that $\beta_{t+1} \notin K$, then there exists a projection facility that pushes β_{t+1} back into the interior of K . Along the boundary of K , the gradient vector of (3.10) is pointing to interior of K .
- (2) $\bar{\Psi}(\beta)$ is a continuously differentiable function, and $\exists \beta^e$ in the interior of K such that $\bar{\Psi}(\beta^e) = 0$. β^e is a locally stable in the sense of Lyapunov.
- (3) Given $\beta \in \mathbb{R}^m$, the probability distribution induced by (3.8) is i.i.d over time with a full support over \mathbf{X} , and the decision maker knows the distribution of ξ_t .
- (4) $\exists M > 0$ such that

$$(3.11) \quad \mathbb{E} [\exp\langle \alpha, \Psi(\beta, X, Z) \rangle] \leq \exp |\alpha| M \quad \forall \alpha \in \mathbb{R}^m.$$

Under Assumption 3.1, the set of sample paths induces by (3.9) converges to the trajectory of (3.10) weakly, as $a \rightarrow 0$ and $t \rightarrow \infty$: $\exists \epsilon' > 0$ such that $\forall \beta_1 \in \mathcal{N}_{\epsilon'}(\beta^e)$, $\forall \epsilon \in (0, \epsilon')$,

$$\lim_{t \rightarrow \infty} \lim_{a \rightarrow 0} \mathbf{P} (|\beta_t - \beta^e| \leq \epsilon) = 1.$$

See Dupuis and Kushner [18] for a formal proof. At β^e , the beliefs of the decision maker are confirmed, even if the decision maker's model is not properly specified. Thus, we call β^e the *self-confirming* equilibrium [37].

The third condition implies that the innovation in each period is i.i.d. for a *given* perception of the decision maker. We can easily expand the model to cover the case where the innovation term evolves according to an ARMA process. Note that as β_t changes in response to the sequence of observations, the regression residual need not remain i.i.d. The assumption that the decision maker knows the correct distribution ξ_t differentiates our exercise from the robust decision making problem analyzed in [33, 34]. After presenting the main result, we shall explore the consequences of dropping this assumption so that the decision maker entertains more general forms of model uncertainty.

While we are intentionally avoiding distributional uncertainty, the decision maker is still exposed to two kinds of model misspecification. First, his model may be under-specified, as he is not aware of \mathbf{Z} . Second, and more importantly, his model may not capture the complex interaction between the evolution of his model and the evolution of the state variables. The second misspecification differentiates the analysis of this paper from the econometric literature on specification testing [41], where the data generating mechanism is exogenously given.

Condition (4) is commonly known as Cramer's condition, which essentially implies the existence of the moment generating function of $\Psi(\beta, X, Z)$. This is a standard condition to ensure that the distribution of $\Psi(\beta, X, Z)$ does not have excessively thick tails.

Because we need to investigate the large deviation properties of (3.9), we need additional notation. Define the H -functional as

$$(3.12) \quad H(\alpha, \beta, t) = \lim_{\tau \rightarrow 0} \sup \lim_{a \rightarrow 0} \sup \frac{\tau}{a} \log \mathbb{E} \left[\exp\langle \alpha, \sum_{k=1}^{\lceil \tau/a \rceil} \Psi(\beta_{t+k}, X_{t+k}, Z_{t+k}) \rangle \mid \beta_t = \beta, \mathcal{H}_t \right],$$

where \mathcal{H}_t is the sigma algebra generated by information at t . The Legendre transform of the H -functional is defined as

$$(3.13) \quad L(\beta, \zeta, t) = \sup_{\alpha} [\langle \alpha, \zeta \rangle - H(\alpha, \beta, t)]$$

and the action functional is then defined as

$$(3.14) \quad S(\beta, T, \phi) = \int_0^T L(\dot{\phi}, \phi, t) dt$$

where $\phi(0) = \beta$ and ϕ is absolutely continuous; otherwise, $S(\beta, T, \phi) = \infty$. The action functional captures the ‘cost’, in probabilistic terms, of any given path. Less likely paths are assigned higher costs.

Dupuis and Kushner [18] show that if H is α -differentiable at $\alpha = 0$ then (3.9) satisfies the large deviation upper bound:¹⁴

$$(3.15) \quad \lim_{a \rightarrow 0} a \log \mathbf{P}(\beta \in A \mid \beta(0) = \beta^e) \leq -\sup_t \inf_{A_t} S(\phi(0), t, \phi) = -S^* < 0$$

where

$$A_t = \{\phi \mid \phi(0) = \beta^e, \exists t, \phi(t) \notin N_\delta(\beta^e)\}$$

and the equality holds along the dominant escape path out of $N_\delta(\beta^e)$. In general, calculating the large deviation rate function, S^* , is a challenging calculus of variations problem. However, Williams [42] shows that in Linear-Quadratic/Gaussian settings the action functional simplifies to the following quadratic form:

$$(3.16) \quad S(\beta, T, \phi) = \int_0^T [\dot{\phi} - \bar{\Psi}(\phi)]' Q^\dagger [\dot{\phi} - \bar{\Psi}(\phi)] dt \quad \phi(0) = \beta$$

where the weighting matrix Q contains information on the likelihood, or ‘cost’, of departures from the mean dynamics. For example, at a given value of β , the instantaneous escape direction is just given by the eigenvector associated with the smallest eigenvalue of Q^\dagger . The full dynamic path can be obtained by solving a set of matrix Lyapunov equations.

3.2. Validation Dynamics. Consider a decision maker who is aware of possible model misspecification, and therefore runs a specification test before using the model to guide his decision. Because the objective of this process is in a certain sense to validate his current model, we call the induced dynamics of the regression coefficients and the observed state variables the *validation dynamics*. This alternative behavior is an abstraction of the casual observation that policy makers search for better models by routinely running specification tests and selecting a best possible model for a given set of available data. Instead of focusing on a specific testing method, we shall consider a class of rules that satisfy fairly common properties.

As in the recursive learning models, $t = 1, 2, \dots$ represents ‘calendar’ time. We call the model the decision maker uses to guide policy a *reference model*. Let $\gamma_k \in \mathbb{R}^m$ be the k -th reference model. Because the same reference model can be used over many periods, typically $k \leq t$.

¹⁴By duality, the α -differentiability condition can be replaced by the condition that there exists a unique ζ^* satisfying $L(\beta, \zeta^*, t) = 0$.

We assume that the reference model has the same m parameters as the recursive learning model, because the reference model is also subject to possible misspecification. If γ_k is validated, the decision maker chooses an action and $y_t = b^r(\gamma_k, X_t, Z_t) \in \mathbf{X}$ is realized in period t .

To be validated, the model parametrized by γ_k must pass a relative entropy specification test. To illustrate the validation process, it is more convenient to represent a model by a probability distribution on \mathbf{X} induced by γ_k . Because of the feedback feature, γ_k influences the underlying probability distribution, in particular, of $X \in \mathbf{X}$. Let \mathcal{M}_{γ_k} be the probability distribution induced by γ_k . The causality here is very important. Unless γ_k is a self-confirming equilibrium, γ_k need not be the best fit model for the distribution \mathcal{M}_{γ_k} .

The test statistic is constructed from the empirical distribution \mathcal{M}_t :

$$(3.17) \quad \mathcal{M}_t(A) = (1 - a) \sum_{\ell=t_k}^t a^{t-\ell+1} \mathbf{1}_{X_\ell \in A}$$

where t_k is the first point in time when the present reference model γ_k is implemented. If the reference model is replaced, then the decision maker should discard old data, because they are generated under different regimes. However, even if one fixes $t_k = 1 \forall k$, the same analysis goes through.

We use relative entropy to define the discrepancy between the two probability distributions, which requires the existence of the Radon-Nikodym derivative. To this end, we need to “smooth” out the empirical distribution so that the test statistic has full support over \mathbf{X} . Let $\hat{\mathcal{M}}_t$ be the smoothed empirical distribution: $\hat{\mathcal{M}}_t$ is atomless and the support of $\hat{\mathcal{M}}_t$ is \mathbf{X} and

$$(3.18) \quad \hat{\mathcal{M}}_t - \mathcal{M}_t \rightarrow 0$$

with probability 1 as $t \rightarrow \infty$. We intentionally remain vague about the specific details of the smoothing method which might depend on the characteristics of the assumed X distribution. For example, in the Gaussian case it suffices to estimate recursively the first two moments from the data to build $\hat{\mathcal{M}}_t$.

Define

$$I_{t,k}^a = I(\hat{\mathcal{M}}_t \| \mathcal{M}_{\gamma_k}) = \int \log \frac{d\hat{\mathcal{M}}_t}{d\mathcal{M}_{\gamma_k}} d\hat{\mathcal{M}}_t$$

whenever $\frac{d\hat{\mathcal{M}}_t}{d\mathcal{M}_{\gamma_k}}$ is well defined. Otherwise,

$$I(\hat{\mathcal{M}}_t \| \mathcal{M}_{\gamma_k}) = \infty.$$

as the relative entropy between the smoothed empirical distribution and the reference model. Clearly, (3.18) implies that $I_{t,k}^a$ is a consistent estimator of the relative entropy: $\forall a, k, I_{t,k}^a - I(\mathcal{M}_t \| \mathcal{M}_{\gamma_k}) \rightarrow 0$ with probability 1.

Because the empirical distribution is constructed recursively after a new reference model is implemented, it is reasonable to assume that $I_{t,k}^a$ can be approximated by a deterministic process.

Assumption 3.2. As $a \rightarrow 0$, $\{I_{t,k}^a\}$ can be approximated by a trajectory of the deterministic path induced by

$$\dot{I}_{t,k} = \varphi(I_{t,k})$$

where φ is continuous. Let t_k be the first period when the k -th reference model is adopted. Then, $\forall k, \forall \tau > 0$,

$$\lim_{a \rightarrow 0} \mathbb{E} \left[\left(\sum_{t=t_k}^{t_k + \lceil \tau/a \rceil} I_{t,k}^a - \int_0^\tau \varphi(I_{s,k}) ds \right)^2 \mid I_{t_k,k}^a = I_{0,k} \right] = 0$$

We can now describe formally the validation process. Let $\gamma_1 = \gamma \in \mathbb{R}^m$ be the initial reference model, and $\rho_1 = 0$ be the initial reference value. Fix $\rho > 0$, $k \geq 1$ and $t \geq 1$. Given reference model γ_k in period t , test statistics $I_{t,k}^a$, a reference value ρ_k , and threshold $\rho > 0$, the decision maker uses γ_k to choose an action in period t if and only if

$$(3.19) \quad I_{t,k}^a \leq \rho + \rho_k.$$

If $I_{t,k}^a > \rho + \rho_k$, the decision maker discards γ_k and choose γ_{k+1} by solving

$$(3.20) \quad \rho_{k+1} = \min_{\gamma} I(\hat{\mathcal{M}}_{t,k}^a \parallel \mathcal{M}_{\gamma})$$

and if there exist multiple solutions, the decision maker chooses the one with the fewest parameters.

$\mathcal{M}_{t,k}^a$ is absolutely continuous with respect to \mathcal{M}_{γ} . The minimized value provides the reference value ρ_{k+1} for γ_{k+1} , which is in fact the lowest relative entropy possible. Essentially, the decision maker is constantly searching for a “better” model in the sense of reducing relative entropy.

If $m = n$ and the model is properly specified in a statistical sense, then (3.19) is just a likelihood ratio test, and (3.20) describes maximum likelihood estimation. In contrast to standard statistical testing problems, the notion of a “true model” is changing over time here. Thus, it is not unusual for the decision maker to reject a misspecified model and replace it by another misspecified model. Thus, our selection process is close to what Vuong [41] proposes. Yet, the alternative hypothesis is not well specified in our case, while the null hypothesis is that the reference model is an accurate description of the state. In this sense, our testing procedure is the non-robust version of the universal testing method [43, 33, 34].

In case of the Gaussian perturbations, as illustrated in section 2, $\hat{\mathcal{M}}_t$ is induced by a maximum likelihood estimator. Thus, once the reference model is rejected, the decision maker can borrow the maximum likelihood estimator from $\hat{\mathcal{M}}_t$ to construct γ_{k+1} . In this case, $\rho_{k+1} = 0$. In general, if we restrict $m \ll n$, then ρ_k is typically bounded away from 0. In a certain sense, $\rho_k > 0$ is the penalty the decision maker pays for using a restricted class of models.

For later reference, let us define the H -functional for the validation dynamics:

$$(3.21) \quad H^{\rho}(\alpha, \gamma, t) = \lim_{\tau \rightarrow 0} \sup \lim_{a \rightarrow 0} \sup \frac{\tau}{a} \log \mathbb{E} \left[\exp \langle \alpha, \sum_{k'=0}^{\lceil \tau/a \rceil - 1} \Psi(\gamma_k, X_{t+k'}, Z_{t+k'}) \rangle \mid \gamma_t = \gamma, \mathcal{H}_t \right],$$

and similarly, we can define the mean dynamics for the validation dynamics as

$$\dot{\gamma} = \bar{\Psi}^\rho(\gamma, x, z)$$

where

$$\bar{\Psi}^\rho(\gamma, x, z) = \lim_{T \rightarrow \infty} \frac{1}{T} \mathbf{E} \left[\sum_{t=1}^T \Psi(\gamma_k, X_t, Z_t) \mid \gamma_1 = \gamma, X_1 = x, Z_1 = z \right]$$

where γ_k evolves according to (3.19) and (3.20).

4. RESULTS

We shall now prove that the sample path properties of the validation dynamics converge to those of the learning dynamics as the testing threshold goes to zero. As pointed out in section 2, it is not enough to prove that the trajectories of the mean dynamics are close. We must also show that the large deviation properties get close as $\rho \rightarrow 0$.

It is instructive to first examine the convergence of the mean dynamics. This implies that if the initial model is not a self-confirming equilibrium, the recursive learning dynamics and the validation dynamics push the model along the same trajectory.

Proposition 4.1. $\forall \gamma$ where $\bar{\Psi}(\gamma) \neq 0$.

$$\lim_{\rho \rightarrow 0} \bar{\Psi}(\gamma) - \Psi^\rho(\gamma) = 0.$$

Proof. The basic idea of the proof is simple. If γ is not a self-confirming equilibrium, then the reference model is rejected in finite time with probability one, because the decision maker's action changes the underlying probability distribution. For small $a > 0$, the interval that a reference model survives converges to a deterministic time. That time interval shrinks as the decision maker runs a tighter specification test: $\rho \rightarrow 0$. Thus, the reference model is updated frequently as in the recursive learning model.

To formalize this intuition, recall that the evolution of the test statistic $I_{t,k}^a$ converges to a deterministic process

$$\dot{I}_{t,k} = \varphi(I_{t,k}).$$

$\forall I_{t,k}$, it takes $\tau_{k,\rho}$ amount of time to get rejected.

Fix $\tau > 0$ and consider

$$\mathbf{E} \sum_{t=1}^{\lceil \tau/a \rceil} a \Psi(\gamma_k, X_t, Z_t)$$

with $\gamma_1 = \gamma$. Let T_k^a be the random time when the k -th reference model is rejected. Since the test statistics converges to a deterministic process, we know $\exists \zeta_k > 0$ such that

$$a(T_{k+1}^a - T_k^a) \rightarrow \zeta_k$$

as $a \rightarrow 0$. Thus,

$$\begin{aligned} \mathbb{E} a \sum_{t=1}^{\lceil \tau/a \rceil} \Psi(\gamma_k, X_t, Z_t) &= \mathbb{E} \sum_{k=1}^K a \sum_{t=T_k^a+1}^{T_{k+1}^a} \Psi(\gamma_k, X_t, Z_t) \\ &= \mathbb{E} \left[\sum_{k=1}^K a(T_{k+1}^a - T_k^a) \left[\mathbb{E}_{T_k^a} \sum_{t=T_k^a+1}^{T_{k+1}^a} \frac{1}{T_{k+1}^a - T_k^a} \Psi(\gamma_k, X_t, Z_t) \right] \right]. \end{aligned}$$

As $a \rightarrow 0$,

$$\mathbb{E}_{T_k^a} \sum_{t=T_k^a+1}^{T_{k+1}^a} \frac{1}{T_{k+1}^a - T_k^a} \Psi(\gamma_k, X_t, Z_t) \rightarrow \Psi(\gamma_k, X_{t_k}, Z_{t_k})$$

and

$$a(T_{k+1}^a - T_k^a) \rightarrow \zeta_k$$

with probability 1. Thus,

$$\mathbb{E} \sum_{t=1}^{\lceil \tau/a \rceil} a \Psi(\gamma_k, X_t, Z_t) \rightarrow \mathbb{E} \sum_{k=1}^K \zeta_k \Psi(\gamma_k, X_{t_k}, Z_{t_k}).$$

As $\rho \rightarrow 0$, the right hand side is an average of $\Psi(\gamma_k, X_{t_k}, Z_{t_k})$. As $\tau \rightarrow 0$, we obtain (3.10). Q.E.D.

However, this logic does not apply if γ is the self-confirming equilibrium, because the mean dynamics around the self-confirming equilibrium vanish. Around the self-confirming equilibrium, we need to compare probability distributions of unlikely events, which are governed by large deviation properties. While the analysis of large deviations in principle requires us to minimize the action functional, Kushner [27] proves that if two H -functionals converge uniformly, then so do the large deviation properties, such as escape paths and expected escape times.

Theorem 4.2. *If $\bar{\Psi}(\gamma) = 0$, then*

$$\lim_{\rho \rightarrow 0} H^\rho(\alpha, \gamma, t) - H(\alpha, \gamma, t) = 0 \quad \forall \alpha, \forall t.$$

Proof. Note that

$$H(\alpha, \gamma, t) = \log \mathbb{E} \exp\langle \alpha, \Psi(\gamma, X, Z) \rangle.$$

Fix $\tau > 0$. Because $\bar{\Psi}(\gamma) = 0$, the reference model is the self-confirming equilibrium, and therefore, rejection is an unlikely event when $a \rightarrow 0$. By (3.15), characterize the bound for the probability: $\forall \tau > 0$,

$$\mathbb{P}(\forall T_k^a, aT_k^a \geq \tau \mid \gamma = \beta^e) \geq 1 - e^{-S^*/a}.$$

Because the rejection time is random, it is necessary to examine two separate cases.

CASE 1. $\forall T_k^a, aT_k^a \geq \tau$

In this case, which occurs with probability at least $1 - e^{-S^*/a}$, the initial reference model is *not* rejected. Thus, throughout $\lceil \tau/a \rceil$ periods, the decision maker continues to take the same action. Thus, the H functional in this case is

$$\begin{aligned} & \lim_{\tau \rightarrow 0} \lim_{a \rightarrow 0} \log \mathbb{E} \left[\exp \left\langle \alpha, \sum_{k'=1}^{\lceil \tau/a \rceil} \Psi(\gamma, X_{t+k'}, Z_{t+k'}) \right\rangle \mid \mathcal{H}_t, \forall T_k^a, aT_k^a \geq \tau \right] \\ &= \log \mathbb{E} \exp \langle \alpha, \Psi(\gamma, X, Z) \rangle = H(\alpha, \gamma, t). \end{aligned}$$

CASE 2. $\exists T_k^a : aT_k^a < \tau$

This is a small probability event, but also is the case where γ_k is changing. As a result, we have little idea about the size of the forecasting error, because γ_k is selected to minimize the Kullback-Leibler distance rather than the forecasting error. Here, we have to use the rate function in combination with the fact that we can choose $\tau > 0$ arbitrarily small before letting $a \rightarrow \infty$. Although we have little idea about the forecasting error, (3.11) in Assumption 3.1 implies that $\exists M > 0$ such that

$$\mathbb{E} [\exp \langle \alpha, \Psi(\beta, X, Z) \rangle] \leq \exp |\alpha| M.$$

Thus,

$$\begin{aligned} & \lim_{\tau \rightarrow 0} \lim_{a \rightarrow 0} \log \mathbb{E} \left[\exp \left\langle \alpha, \sum_{k'=1}^{\lceil \tau/a \rceil} \Psi(\gamma, X_{t+k'}, Z_{t+k'}) \right\rangle \mid \gamma_1 = \gamma, \mathcal{H}_t, (\exists T_k^a : aT_k^a < \tau) \right] \mathbb{P} (\exists T_k^a : aT_k^a < \tau) \\ & \leq \lim_{\tau \rightarrow 0} \lim_{a \rightarrow 0} \log \left[\exp \left(\frac{|\alpha| M (\tau + a) - S^*}{a} \right) \right]. \end{aligned}$$

where the inequality uses both the first part of Assumption 3.1 and the large deviations rate function in (3.15). Choose $\tau > 0$ sufficiently small that

$$\tau < \frac{S^*}{2|\alpha|M}.$$

Then,

$$\limsup_{a \rightarrow 0} |\alpha| M (\tau + a) - S^* < 0.$$

Thus, the right hand side converges to 0 as $a \rightarrow 0$.

Combining the two cases, we conclude that the first case dominates in determining the value of $H^\rho(\alpha, \gamma, t)$, which is precisely $H(\alpha, \gamma, t)$. Thus, the H -functional of the validation dynamics converges uniformly to that of the recursive learning dynamics. *Q.E.D.*

5. DISCUSSION AND EXTENSIONS

The previous analysis is quite general, and can be extended in a number of directions. In this section, we briefly discuss two that are likely to be important in applications. First, we show that lagged variables and two-sided learning can be incorporated into the analysis. Second, we comment on how the recent work of Meyn et al on robust hypothesis testing could be used to relax our assumption that agents know the true error distribution. We also speculate that this might lead to an alternative method of calibrating robustness.

5.1. Lagged variables and two-sided learning. In a companion paper (Cho and Kasa [8]), we argue that model validation dynamics might be a contributing factor in observed financial crises. The analytical framework is similar to the one used here, but in order to fit some important features of observed crises, we had to relax two assumptions. First, in order to generate real effects from crises, we needed to relax the ‘Fed watcher’ assumption. In this case, both the government and the private sector must learn. Second, to match observed persistence in the real effects, we needed to add a lag in the output equation. We use the well known third-generation crisis model of Aghion, Bacchetta, and Banerjee [1] to do this.

Consider a policy maker who wants to control the exchange rate, s_t , and output, y_t , to achieve some objective. This objective could reflect some well defined notion of social welfare, or alternatively, as we assume, be an ad hoc quadratic loss function:

$$(5.22) \quad V = \min_{\{s_t\}} E_0^g \sum_{t=0}^{\infty} \delta^t [\lambda(y_t - y^*)^2 + (s_t - s^*)^2]$$

where E^g denotes expectations based on the government’s beliefs. Output and the exchange rate are related to each other by the following dynamic ‘expectations-augmented Phillips curve’:

$$(5.23) \quad y_t = \bar{y}(1 - \alpha) + \alpha y_{t-1} + \theta(s_t - E_{t-1}^p s_t) + \sigma_1 \varepsilon_{1t} \quad \varepsilon_{1t} \sim N(0, 1)$$

where E^p denotes expectations based on the private sector’s beliefs. Note that with Rational Expectations (and common information sets), $E^p = E^g$, and they both use the true objective probability distribution. Here we do not impose the Rational Expectations Hypothesis. Instead, we assume the policy maker must learn about the relationship between the exchange rate and output by witnessing the response of output to his sequence of exchange rate choices. Likewise, the private sector must learn about the government’s exchange rate policy, using the observed histories of output and the exchange rate and a guess about the functional form of the government’s policy function.¹⁵ In contrast to Bayesian learning, we do not allow for any experimentation or strategic interaction in these learning problems. Both sides learn, but purely in a passive, retrospective way.¹⁶

The government’s approximating model is assumed to take the following form:

$$(5.24) \quad y_t = \beta_0 + \beta_1 y_{t-1} + \beta_2 s_t + u_{1t}$$

As before, this imputes a subtle form of specification error to the government. Whereas in reality it is only unanticipated devaluations that matter, the government mistakenly believes that the exchange rate by itself matters. As a result, the evolving beliefs of the private sector inject ‘parameter drift’ into the government’s approximating model.

¹⁵In [8] we show that under certain conditions the large deviations properties of the model are governed solely by the beliefs of the government, even when the private sector’s beliefs are adaptive.

¹⁶Kreps [26] provides arguments in favor of this approach to learning. He calls it ‘anticipated utility’. Cogley and Sargent [14] show in the context of a standard permanent income model that anticipated utility and Bayesian learning produce very similar outcomes. However, they also caution that the two learning strategies can depart significantly when agents are highly risk averse.

The government solves the optimization problem in (5.22) given its perceived model in (5.24). This produces the following policy function

$$(5.25) \quad s_t^p = g_0(\beta) + g_1(\beta)y_{t-1}$$

where $\beta = (\beta_0, \beta_1, \beta_2)$, and $g_0(\cdot)$ and $g_1(\cdot)$ are differentiable functions of β . Following Sargent and CWS, assume that the actual exchange rate is equal to the planned exchange rate, s_t^p , plus an i.i.d. shock, which captures random implementation errors or high frequency money demand shocks. Thus, the market exchange rate is,

$$(5.26) \quad s_t = s_t^p + \sigma_2 \varepsilon_{2t} \quad \varepsilon_2 \sim N(0, 1)$$

In this model the only action the private sector takes is to forecast the exchange rate. It does this using an econometric model. Assume that its model is correctly specified, in the sense that it is consistent with actual government behavior:

$$(5.27) \quad s_t = \gamma_0 + \gamma_1 y_{t-1} + u_{2t}$$

Note that with CWS's Fed watcher assumption we would have $\gamma_0 = g_0(\beta)$ and $\gamma_1 = g_1(\beta)$.

Substituting (5.25)-(5.27) into (5.23) then delivers the *actual* law of motion as a function of the beliefs of the government and private sector:

$$(5.28) \quad y_t = \bar{y}(1 - \alpha) - \theta\gamma_0 + (\alpha - \theta\gamma_1)y_{t-1} + \theta s_t + \sigma_1 \varepsilon_{1t}$$

where s_t is given by (5.26). Because of model misspecification, agents' beliefs will not converge to a Rational Expectations Equilibrium. Instead, they converge to a *self-confirming equilibrium*, defined to be a situation where agents no longer have an incentive to revise their models. In [8] we show this model has a unique self-confirming equilibrium, and that it is E-stable if $0 < \alpha < 1$.

As before, we assume agents update their beliefs using constant-gain stochastic approximation algorithms. The only difference is that now, with two-sided learning, we must specify *two* gain parameters. While it might be interesting to explore the effects of differential learning rates, for most of our analysis we assume the gain parameter of the government equals the gain parameter of the private sector.

Figure 3 reports representative sample paths from this model. The key parameter values are as follows: (1) $a_p = a_g = .04$, which implies a half-life of data relevance of about 17 time periods, (2) $\theta = -0.3$, which reflects our assumption that adverse balance sheet effects dominate liquidity effects, so that unanticipated devaluations are *contractionary* (see [8] for details), (3) $\lambda = 1.5$, which implies that output fluctuations are more costly than exchange rate fluctuations, (4) $\bar{y} = 1.0$ and $y^* = 1.2$, which reflects the usual assumption in these models that the target output level exceeds the natural rate, and (5) $\sigma_1^2 = .0003$ and $\sigma_2^2 = .0001$, which implies that real shocks are more volatile than nominal shocks.

Three main features jump out at you in these simulations. First, at least in a qualitative sense, the exchange rate paths resemble the observed exchange rate histories of many crisis prone countries. There are prolonged periods of gradual appreciation, followed by rare but *recurrent* crisis episodes, where the exchange rate depreciates sharply. In this particular case, they occur about once every 1000 periods, i.e., about once every 4-5 years if the time unit is a day, or about once every 20 years if the time unit is a week. Increasing the gain parameters or the shock variances increases the frequency of crises. Second, *crises cause recessions*, with output typically falling by about 10% during a crisis. Of course, this is

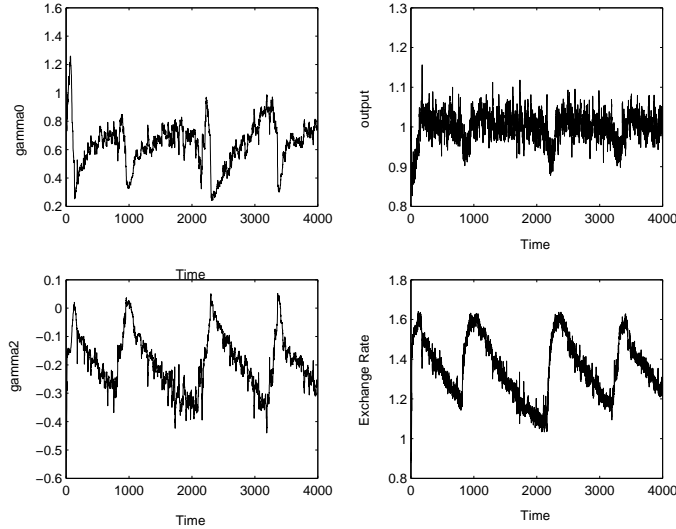


FIGURE 3. Example Simulation of Dynamic Model: $a_g = a_p = .04$

not a coincidence. Generating contractionary devaluations is the *raison d'être* of the third-generation crisis literature. Here it is entirely driven by our assumption that θ is negative. Finally, the third feature to notice is that crises are intimately connected to the evolving beliefs of the government. The two plots in the left-hand side of Figure 3 depict the paths of the coefficient estimates. The up-and-down pattern of the exchange rate coefficient estimate mirrors the up-and-down pattern of the exchange rate. What's happening here is that the government is vacillating between Sargent's [35] two observationally equivalent ways of interpreting the data. Crises occur when the government confuses the natural rate properties of the model with the apparent absence of balance sheet effects. In [8] we provide a more detailed account of these dynamics.

For our purposes here, it is more useful to illustrate how the validation dynamics of this model converge to the recursive learning dynamics. To do this, we simulated the model 30 times (each one consisting of $T = 3000$ time periods) for two different values of the relative entropy threshold, ρ . Figures 4 and 5 depict a representative outcome when $\rho = .08$. The top row of Figure 4 reports values of relative entropy and a $(0, 1)$ indicator variable for model rejection. Vertical lines represent a rejection. The middle row shows the paths of output and the exchange rate for the validation dynamics, and the bottom row does the same thing for the learning dynamics.

For this particular run, the average time between model rejections is 34 periods. However, notice that during an escape rejections arrive more frequently. This can also be seen in Figure 5, which reports the paths of the coefficient estimates for the same simulation. The top row contains the learning dynamics and the bottom row contains the validation dynamics.

Notice that while the model is converging to the self-confirming equilibrium there can be prolonged periods without model revision. However, once a self-confirming equilibrium

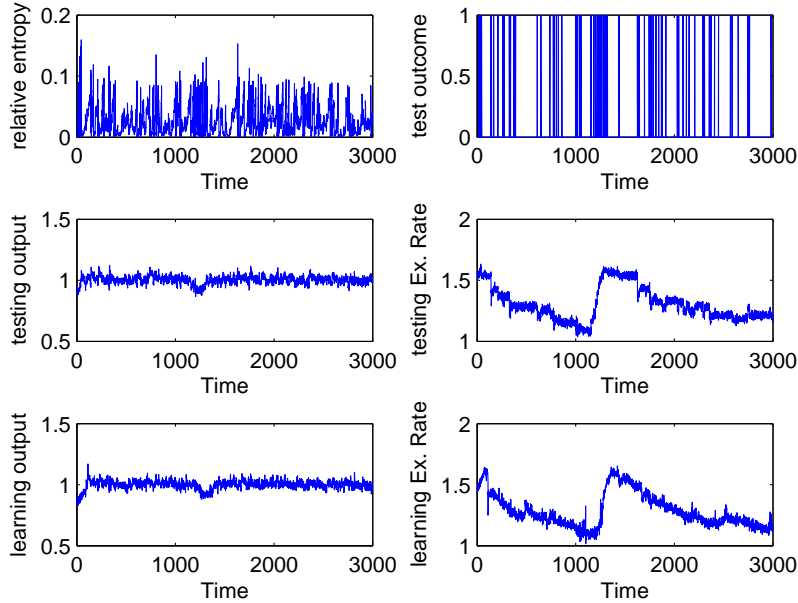


FIGURE 4. Representative Sample Paths of the Validation and Learning Dynamics: $\rho = .08$

is rejected, and an escape is ignited, model rejections pile up and the coefficient estimates move rapidly toward their (distorted) natural rate values.

Although Figures 4 and 5 are interesting, they provide no information on the convergence of the validation dynamics to the learning dynamics. For that we need to compare simulations for alternative values of ρ . Also, we need to recognize that there is going to be some sampling variability. Larger ρ values can sometimes produce closer alignment between validation and testing dynamics. We need to make sure we don't get unlucky (or lucky, for that matter). Hence, Figs 6 and 7 compare the *distributions* of outcomes for different values of testing threshold ρ .

In particular, for each of the 30 simulations we computed the L^2 distance between the exchange rate paths induced by the model validation dynamics and the recursive learning dynamics. That is, if $s^{mv}(t)$ denotes the model validation path and $s^{rl}(t)$ denotes the recursive learning path, then we compute the discrete analog of:

$$(5.29) \quad D = \left(\int_0^{3000} (s^{mv}(t) - s^{rl}(t))^2 dt \right)^{1/2}$$

Note that this is actually a *weaker* sense of convergence than our theorem predicts, since it does not restrict tail discrepancies. Nonetheless, it does give some sense of sample path convergence.

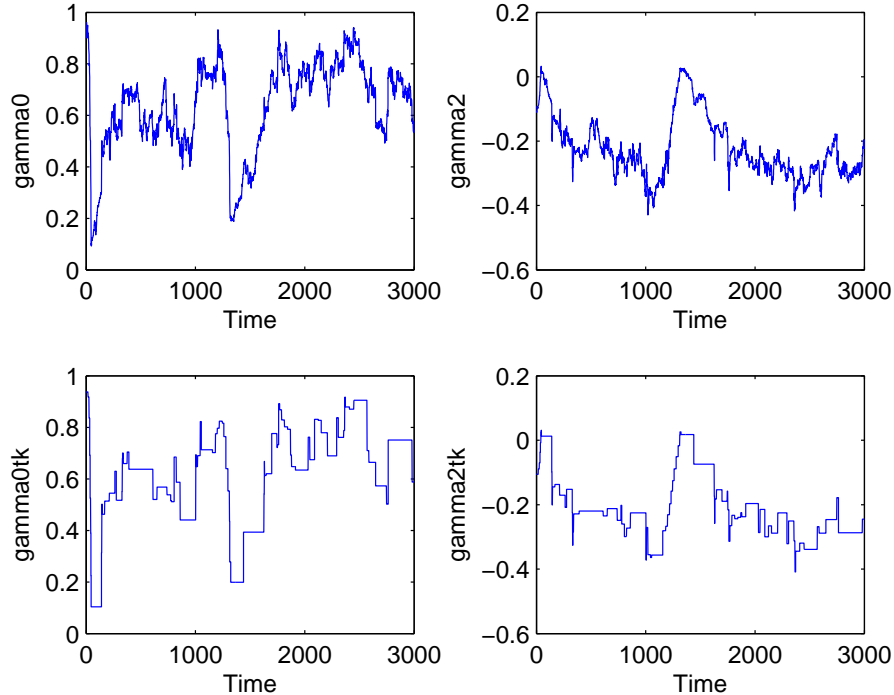


FIGURE 5. Sample Paths of the Coefficient Estimates: $\rho = .08$

Figure 6 reports smoothed histograms for $\rho = .03$ and $\rho = .08$.¹⁷ Although not dramatic, Figure 6 depicts a clear sense of convergence. A smaller value of ρ causes the exchange rate paths to become more aligned and the distribution of L^2 distances to shift left.

What lies behind this convergence is a more rapid rate of model revision. Figure 7 reports model rejection rates for the same set of simulations. For each of the 30 runs we computed the average time to rejection by dividing 3000 by the total number of rejections. As discussed earlier, this is a bit misleading, since rejection rates are *not* constant over time, i.e., they occur more frequently during escapes and less frequently while the model is in the neighborhood of the self-confirming equilibrium. Once again, we smoothed the raw histograms with a kernel density smoother.

Figure 7 shows that when $\rho = .08$ the average reference model lasts for about 40 periods before it is rejected, i.e., a little bit less than a year if the time unit is a week. When ρ falls to .03 the average survival time drops to only about 25 periods.

5.2. Robust validation. Earlier we alluded to a debate between Sims [40] and Cogley and Sargent [12] about the presence of regime changes in U.S. inflation data. Sims points out that if agents fit models with homoskedastic error terms when in fact the data are heteroskedastic, they may be fooled into inappropriately inferring that there have been

¹⁷The histograms were smoothed with matlab routine `KSDENSITY`, using the default gaussian kernel.

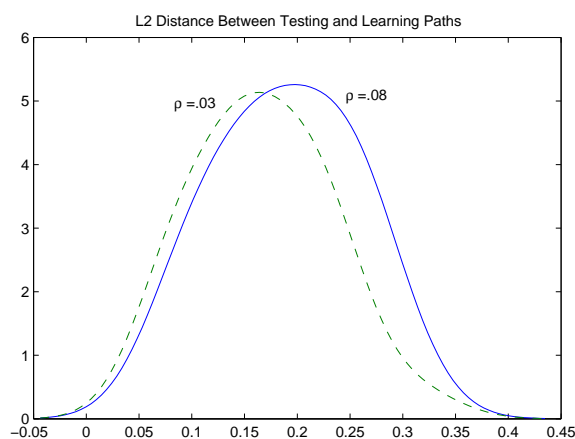


FIGURE 6. Probability Distribution of the L^2 Distance between Validation and Learning Exchange Rate Paths

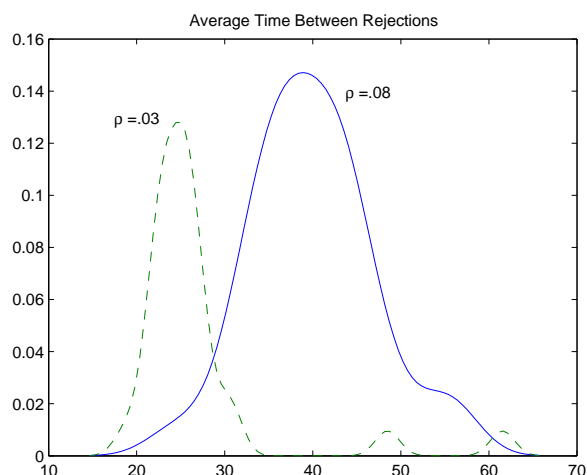


FIGURE 7. Probability Distribution of the Average Time Until Model Rejection

breaks in the data, and as a result, inappropriately reject their models. Sims' point is relevant for us too. Until now we have assumed the decision maker knows the model's error distribution, even if he doesn't know its parameters. What if his beliefs about this distribution are wrong? If he ignores this possibility then he exposes himself to the kind of error that Sims highlighted. Our goal here is not just to model policymakers as econometricians, but to model them as *good* econometricians. Good econometricians worry about *robustness*.

If the only thing he had to worry about was parameter estimation and distributional uncertainty, there would be a straightforward response - use GMM rather than maximum

likelihood. However, the agent we model is not just an econometrician, he is a decision maker who *uses* his model to devise a control policy. Moreover, this control policy influences the data-generating process.¹⁸ This means that he must make enough assumptions about his environment that he can solve his control problem, and this necessarily exposes him to greater specification risk than if he just needed to estimate parameters.

Our approach to this problem is to blend the recent literature on robust control and filtering (Hansen and Sargent [22]) with the recent literature on robust inference in moment condition models (Kitamura and Stutzer [25] and Kitamura and Otsu [24]). We do this by building on the work of Pandit [33] and Pandit and Meyn [34]. As in the econometric literature on information-theoretic GMM and empirical likelihood, we define models by parameterized *moment conditions*. However, for us, these moment conditions do not come from economic theory; they define a permitted class of *model perturbations*, within which Hansen and Sargent’s ‘evil agent’ can select a model to subvert the agent’s model validation and control efforts. The more moment conditions there are, the less freedom the evil agent has, and hence, the less robust will be the outcome.

An important by-product of a robust model validation approach is that it endogenously delivers a ‘robustified’ reference model. Existing work on robust control is silent about where the robustness-seeking agent’s reference model comes from. It only considers perturbations to a *given* reference model. In addition, due to the links between KLIC and Type I and Type II error rates, our approach endogenously generates *detection errors*. Current work on robustness specifies these errors exogenously, as a device to calibrate ‘reasonable’ amounts of robustness. (See Anderson, Hansen, and Sargent [2]). Instead, our approach exogenously specifies a set of moment conditions.

So let’s now drop the assumption that the decision maker knows the exact distribution of ξ_t so that the decision maker faces some form of model uncertainty. In this case, it is natural for the decision maker to pursue some form of robustness in the validation and decision making process. We formulate the robust validation process following the framework of [33, 34]. It will be more convenient to represent the regression equation in the form of (3.8). Let ϕ be the marginal distribution of X . If the decision maker knows the distribution of ξ_t in (3.8), he can calculate the probability distribution

$$\sum_{\ell=0}^{\bar{\ell}} X_{t-\ell} \beta_{t-\ell}.$$

However, we now assume the decision maker has only partial information about the distribution. Instead of the exact distribution, the decision maker knows only a few moments. For example, consider the following moment-constrained set of parameterized models

$$\mathbb{P}(\beta) = \left\{ \phi \mid \mathbf{E}^\phi X\beta = 0, \text{ and } \mathbf{E}^\phi (X\beta)^2 = \sigma_\xi^2 \right\}$$

where $\sigma_\xi^2 = \mathbf{E}\xi_t^2$. In this case, the decision maker does not know the exact distribution of the regression disturbance. He only knows the first two moments. The number of restrictions imposed on the moment class can be interpreted either as an expression of the decision maker’s bounded rationality or as an expression of his preference for robustness.

¹⁸Remember, however, that in keeping with our assumption of bounded rationality, we assume the agent ignores this feedback.

If he knows the correct distribution of ξ , he must know every moment of ξ , and therefore, the moment class is subject to an infinite number of constraints. The finiteness of the constraints can be interpreted as a bound on his capacity to process information, which exposes him to model uncertainty.

For $\rho_1 > 0$ and a probability distribution ϕ on \mathbf{X} , define

$$\mathcal{Q}_{\rho_1}(\phi) = \{\phi' \mid I(\phi' \parallel \phi) < \rho_1\}$$

and

$$\mathcal{Q}_{\rho_1}(\mathbb{P}) = \bigcup_{\phi \in \mathbb{P}} \mathcal{Q}_{\rho_1}(\phi).$$

For a small $\rho_1 > 0$, we can interpret $\mathcal{Q}_{\rho_1}(\phi)$ as the class of models which are difficult to differentiate from ϕ . Similarly, $\mathcal{Q}_{\rho_1}(\mathbb{P})$ is the class of models that are difficult to distinguish from models in \mathbb{P} .

The robust validation process can now be defined as follows. Let ϕ_k be the probability distribution over \mathbf{X} induced by the present reference model parameterized by γ_k . Given a ‘smooth’ empirical distribution $\hat{\mathcal{M}}_t$, define

$$\mathcal{L}(\hat{\mathcal{M}}_t, \gamma_k) = \inf_{\phi \in \mathbb{P}(\gamma_k)} I(\hat{\mathcal{M}}_t \parallel \phi)$$

as the ‘worst case’ relative entropy over $\mathbb{P}(\gamma_k)$. If

$$\mathcal{L}(\hat{\mathcal{M}}_t, \gamma_k) < \rho_2,$$

then the decision maker uses γ_k to solve

$$(5.30) \quad \sup_{\mathbf{u}_t} \inf_{\phi \in \mathbb{P}(\gamma_k)} \mathbf{E}^\phi(1 - \delta) \sum_{k'=1}^{\infty} \delta^{k'-1} U(u_{t+k'}, X_{t+k'})$$

where $U(\cdot)$ is the one period payoff, and $\mathbf{u}_t = (u_t, u_{t+1}, \dots)$ is the sequence of controls. If

$$\mathcal{L}(\hat{\mathcal{M}}_t, \gamma_k) \geq \rho_2,$$

then γ_k is discarded, and a new reference model γ_{k+1} is constructed by solving

$$\sup_{\gamma} \inf_{\phi \in \mathbb{P}(\gamma)} I(\hat{\mathcal{M}}_t \parallel \phi).$$

With γ_{k+1} in place of γ_k , the decision maker solves (5.30).

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 [36, 37] on boundedly rational learning dynamics. From econometrics, we’ve borrowed from recent work on robust hypothesis testing [43, 33, 34] and the analysis of misspecified models [41, 21]. 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. (See, e.g., Brock, Durlauf, and West's [5] work on Bayesian model averaging). However, we think there are good reasons to explore alternative approaches. (See, e.g., Hansen and Sargent [22], Kreps [26], and Bray and Kreps [4]).

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 prove that as the government employs an increasingly stringent specification test, the dynamics generated by a process of testing and model revision, which we call *validation dynamics*, converge in a very strong way to the dynamics generated by recursive learning models. This is a useful connection to make, because it enables us to apply the results of Williams [42] and Cho, Williams, and Sargent [10] 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. Perhaps the most promising one is to follow-up on the connections between robust validation, robust control, and robust inference in moment constrained models that were briefly outlined in section 5.2. We are actively pursuing this in ongoing work [9].

REFERENCES

- [1] Phillippe Aghion, Phillippe Bacchetta, and Abhijit Banerjee. Currency Crises and Monetary Policy in an Economy with Credit Constraints. *European Economic Review*, 45:1121–1150, 2001.
- [2] Evan W. Anderson, Lars P. Hansen, and Thomas J. Sargent. A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection. *Journal of the European Economic Association*, 1:68–123, 2003.
- [3] Margaret M. Bray. Learning, Estimation and the Stability of Rational Expectations. *Journal of Economic Theory*, 26:318–339, 1982.
- [4] Margaret M. Bray and David M. Kreps. Rational Learning and Rational Expectations. In George R. Feiwel, editor, *Arrow and the Ascent of Modern Economic Theory*, pages 597–625. New York University Press, 1987.
- [5] William A. Brock, Steven N. Durlauf, and Kenneth D. West. Model Uncertainty and Policy Evaluation: Some Theory and Empirics. University of Wisconsin, 2004.
- [6] William A. Brock and C. Hommes. A Rational Route to Randomness. *Econometrica*, 65:1059–1095, September 1997.
- [7] Xiahong Chen and Halbert White. Convergence of Non-parametric Learning Models. *Journal of Economic Theory*, 82:190–222, 1998.
- [8] In-Koo Cho and Ken Kasa. Model Validation Dynamics and Endogenous Currency Crises. University of Illinois, 2005.
- [9] In-Koo Cho and Kenneth Kasa. Robust Model Validation. University of Illinois, 2005.
- [10] In-Koo Cho, Noah Williams, and Thomas J. Sargent. Escaping Nash Inflation. *Review of Economic Studies*, 69:1–40, 2002.
- [11] James Chu, Maxwell Stinchcombe, and Halbert White. Monitoring Structural Change. *Econometrica*, 64:1045–1065, 1996.
- [12] Timothy Cogley and Thomas J. Sargent. Evolving Post World War II U.S. Inflation Dynamics. In Ben Bernanke, editor, *NBER Macroeconomics Annual, 2001*. MIT Press, 2001.
- [13] Timothy Cogley and Thomas J. Sargent. The Conquest of U.S. Inflation: Learning and Robustness to Model Uncertainty. UC Davis and New York University, 2004.
- [14] Timothy Cogley and Thomas J. Sargent. Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making. UC Davis and New York University, 2005.
- [15] Timothy Cogley and Thomas J. Sargent. Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S. *Review of Economic Dynamics*, 8:262–302, 2005.
- [16] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley, 1991.
- [17] Amir Dembo and Ofer Zeitouni. *Large Deviations Techniques and Applications*. Springer-Verlag, New York, 2nd edition, 1998.
- [18] Paul Dupuis and Harold J. Kushner. Stochastic Approximation and Large Deviations: Upper Bounds and w.p.1 Convergence. *SIAM Journal of Control and Optimization*, 27:1108–1135, 1989.
- [19] George W. Evans and Seppo Honkapohja. *Learning and Expectations in Macroeconomics*. Princeton University Press, 2001.
- [20] Roman Frydman and Edmund S. Phelps. *Individual Forecasting and Aggregate Outcomes*. Cambridge University Press, 1983.
- [21] Lars P. Hansen and Thomas J. Sargent. Seasonality and Approximation Errors in Rational Expectations Models. *Journal of Econometrics*, 55:21–55, 1993.
- [22] Lars P. Hansen and Thomas J. Sargent. *Robustness*. Princeton University Press, 2005.
- [23] Kenneth Kasa. Learning, Large Deviations, and Recurrent Currency Crises. *International Economic Review*, 45:141–173, 2004.
- [24] Yuichi Kitamura and Taisuke Otsu. Minimax Estimation and Testing for Moment Condition Models via Large Deviations. Yale University, 2005.
- [25] Yuichi Kitamura and Michael Stutzer. An Information-Theoretic Alternative to Generalized Method of Moments Estimation. *Econometrica*, 65(4):861–874, July 1997.
- [26] David M. Kreps. Anticipated Utility and Dynamic Choice. In *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983-1997*. Cambridge University Press, 1998.

- [27] Harold J. Kushner. Robustness and Approximation of Escape Times and Large Deviations Estimates for Systems with Small Noise Effects. *SIAM Journal of Applied Mathematics*, 44:160–182, February 1984.
- [28] Tze Leung Lai and J.Z. Shan. Efficient Recursive Algorithms for Detection of Abrupt Changes in Signals and Control Systems. *IEEE Transactions on Automatic Control*, 44:952–966, 1999.
- [29] Robert E. Lucas and Thomas J. Sargent. *Rational Expectations and Econometric Practice*. University of Minnesota Press, 1981.
- [30] Robert E. Lucas, Jr. Econometric Policy Evaluation: A Critique. In Karl Brunner and Alan Meltzer, editors, *The Phillips Curve and Labor Markets*. Carnegie-Rochester Conf. Series on Public Policy, 1976.
- [31] Albert Marcet and Thomas J. Sargent. Convergence of Least Squares Learning Mechanisms in Self Referential Linear Stochastic Models. *Journal of Economic Theory*, 48:337–368, 1989.
- [32] Jacob Marschak. Economic Measurement for Policy and Prediction. In William C. Hood and Tjalling G. Koopmans, editors, *Studies in Econometric Method*, pages 1–26. Wiley, 1953.
- [33] C. Pandit. *Robust Statistical Modeling Based On Moment Classes With Applications to Admission Control, Large Deviations and Hypothesis Testing*. PhD thesis, University of Illinois at Urbana Champaign, University of Illinois, Urbana, IL, USA, 2004.
- [34] C. Pandit and S. P. Meyn. Extremal distributions and worst-case large-deviation bounds. Submitted for publication, 2004.
- [35] Thomas J. Sargent. The Observational Equivalence of Natural and Unnatural Rate Theories of Macroeconomics. *Journal of Political Economy*, 84:631–640, 1976.
- [36] Thomas J. Sargent. *Bounded Rationality in Macroeconomics*. Clarendon Press, 1993.
- [37] Thomas J. Sargent. *The Conquest of American Inflation*. Princeton University Press, 1999.
- [38] Thomas J. Sargent and Noah Williams. Impacts of Priors on Convergence and Escapes from Nash Inflation. *Review of Economic Dynamics*, 8:360–391, 2005.
- [39] Christopher A. Sims. Macroeconomics and Reality. *Econometrica*, 48:1–48, 1980.
- [40] Christopher A. Sims. Comment on Sargent and Cogley’s ‘Evolving Post World War II U.S. Inflation Dynamics’. In Ben Bernanke, editor, *NBER Macroeconomics Annual, 2001*. MIT Press, 2001.
- [41] Quang H. Vuong. Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses. *Econometrica*, 57(2):307–333, March 1989.
- [42] Noah Williams. *Escape Dynamics in Learning Models*. PhD thesis, University of Chicago, 2001.
- [43] Ofer Zeitouni and Michael Gutman. On universal hypotheses testing via large deviations. *IEEE Trans. Inform. Theory*, 37(2):285–290, 1991.

DEPARTMENT OF ECONOMICS, UNIVERSITY OF ILLINOIS, 1206 S. 6TH STREET, CHAMPAIGN, IL 61820
USA

E-mail address: inkoocho@uiuc.edu

URL: <http://www.business.uiuc.edu/inkoocho>

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

E-mail address: kkasa@sfu.ca

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