Macroeconomic Dynamics, **24**, 2020, 1124–1150. Printed in the United States of America.
doi:10.1017/S1365100518000731

PRECAUTIONARY LEARNING AND INFLATIONARY BIASES

CHETAN DAVE
University of Alberta

JAMES FEIGENBAUM
Utah State University

In a canonical monetary policy model in which the central bank learns about underlying fundamentals by estimating the parameters of a Phillips curve, we show that the bank's loss function is asymmetric such that parameter overestimates may be more or less costly than underestimates, creating a precautionary motive in estimation. This motive suggests the use of a more efficient variance-adjusted least-squares estimator for learning about fundamentals. Informed by this "precautionary learning" the central bank sets low inflation targets, and the economy can settle near a Ramsey equilibrium.

Keywords: Asymmetric Least Squares, Adaptive Learning, Time Consistency, Inflationary Biases

1. INTRODUCTION

In conventional macroeconomic policy models, policymakers minimize a symmetric loss function over relevant macroeconomic variables using a model to describe the constraints imposed by the economic environment. In practice, a policymaker often determines optimal policy while simultaneously learning the parameters of her model through an adaptive process. We show that if the policymaker uses the same loss function to evaluate policy results and econometric efficiency, a loss function that is symmetric with respect to macroeconomic observables need not be symmetric with respect to parameter estimates since the policymaker's reaction function is nonlinear in the unknown parameter. This asymmetry induces a precautionary motive that favors a variance-adjusted learning algorithm, use of which we show has important implications for policy dynamics. We therefore argue that one way to model "precaution in policymaking" is via the assumption of a variance-adjusted learning algorithm in frequentist

We thank Jess Benhabib, John Duffy, John Leahy, Tom Sargent, and colleagues and seminar participants at New York University (Abu Dhabi), the Southern Economics Association Conference, Society for Nonlinear Dynamics and Econometrics Conference, the Federal Reserve Banks of Dallas and Richmond, and the Computational Economics and Finance Conference for insightful comments and support. The usual disclaimer applies. Address correspondence to: Chetan Dave, Department of Economics, University of Alberta, 8-14 HM Tory Building, Edmonton, Alberta, Canada T6G 2H4. e-mail: cdave@ualberta.ca.

 1469-8056/18

environments. We are thus able to mimic the intuition of Brainard's (1967) Conservatism Principle.

We adopt the canonical monetary policy context of Kydland and Prescott (1977) and Barro and Gordon (1983) (hereafter, KPBG), in which the central bank minimizes a loss function over inflation and unemployment subject to a Phillips curve.[1] The KPBG model under complete information about the laws of motion of the economy can deliver both a time-inconsistent low inflation Ramsey equilibrium or a time-consistent high inflation Nash equilibrium. Sargent (1999) and Cho et al. (2002) study the impact on targeted inflation when ordinary least-squares regressions are used to learn the parameters of a *misspecified* Phillips curve. Under such adaptive learning the long-run outcome depends on how the policymaker treats new data.[2] For a decreasing-gain learning algorithm in which each new datum is given less weight than past data were at previous iterations, inflation will converge to its Nash equilibrium value. However, for a constant-gain learning equilibrium in which each new datum is weighted the same with each iteration, Cho et al. (2002) find that, under recursive ordinary least-squares (ROLS) learning, inflation will fluctuate near its Nash equilibrium value, but occasionally escape to the Ramsey outcome.

We first show that after expressing inflation and unemployment in terms of the Phillips curve parameters, the loss function is not symmetric since the policymaker's reaction function is nonlinear in the parameter being estimated.[3] As Leland (1968) and Sandmo (1970) established in the context of precautionary saving, an asymmetric loss function will cause optimal decisions to shift under uncertainty. If there is a single stochastic variable in the model, the direction of this shift depends on the sign of the third derivative of the loss function, and the size of the shift is proportional to the variance of the stochastic variable. In our setting, this translates to employing an estimator with a bias proportional to the variance of the unbiased least-squares estimator [Varian (1975) and Zellner (1986)]. We explore the inflation dynamics that result if the central bank employs such "precautionary learning," using a variance-adjusted least-squares (VALS) estimator for the Phillips curve parameters. The VALS estimator augments the ROLS estimator by a term proportional to the variance of the ROLS estimator for each of the two parameters (slope and intercept) of the Phillips curve.[4]

We term the two constants of proportionality the *precautionary parameters*, whose value is zero under the conventional constant-gain ROLS estimator, so the VALS estimator nests the ROLS case. The precautionary parameters measure the strength of the variance correction for the slope and intercept estimators. We will primarily be concerned with the precautionary parameter for the slope since the slope has a bigger impact on the resulting dynamics, but the following discussion would also pertain to the precautionary parameter for the intercept.

The precautionary parameter for the slope can have either sign depending on whether the policymaker is more concerned about overestimating or underestimating the slope because her loss function is higher in that direction. A positive precautionary parameter means the policymaker is adjusting her estimator upward to avoid getting an estimate that is too low, which would be more costly than a high estimate. Conversely, a negative precautionary parameter would be used if she is afraid the least-squares estimate would be too high. Note that the precautionary parameters are parameters of the policymaker's chosen estimator and not parameters of the model. Thus they are choice variables, present in any model of an econometrician, though ordinarily they are, implicitly, set to zero.

We find that the results and intuition underlying escape dynamics are shaped as follows. Under ROLS learning, escapes occur because a sequence of unusual shocks can lead the policymaker to lower her estimate of the Phillips curve slope. This reduces the apparent trade-off between unemployment and inflation, so the policymaker lowers her inflation target, which produces data confirming that the trade-off has indeed diminished and is even smaller than originally thought. The policymaker further reduces the inflation target and the process continues until inflation is driven to zero and the economy replicates the Ramsey outcome. Once the targeted level of inflation is at zero, the negative correlation between inflation and unemployment becomes apparent again. The central bank then ratchets up inflation, and the economy heads back toward the high-inflation Nash outcome.

Given the analytically intractable nature of the model, we simulate the economy and first plot the expected loss as a function of the precautionary parameters. We find that the loss function is not minimized when the precautionary parameters are set to zero. In fact, the ROLS value is near a ridge line of the loss function rather than being close to a minimum. We find that a downward bias of slope estimates increases the frequency of escapes but also increases the long-run level of inflation. The latter effect dominates, increasing the central bank's loss function value. However, we also find that the loss function exhibits a nonmonotonic dependence on the precautionary parameters: for a sufficiently large downward bias, the escape frequency is so high that the economy hardly returns to its high long-run level, and average inflation is very low.[5] In contrast, an upward bias of slope estimates reduces long-run inflation since the central bank believes a small increase in inflation will produce a big decrease in unemployment. It also reduces the frequency of escapes since zero slope estimates are more infrequent. With a large enough bias in either direction, the policymaker achieves average outcomes close to the Ramsey outcome. This can occur with a smaller slope bias if, in addition, the intercept estimate is biased down as well, leading the policymaker to lower her inflation target.[6] As such our result complements that of Brainard (1967) and suggests that the Conservatism Principle may be at work in frequentist environments in which a policymaker adaptively learns about structural parameters.

To help understand how the variance adjustment alters the learning dynamics, we also compute the unconditional distribution of the Phillips curve coefficients and targeted inflation. When the precautionary parameters are set to zero, these distributions are fat-tailed. Escape dynamics are represented by the high probability of loitering for a time in the lower tail of the distribution for targeted inflation. However, if we explore the rest of the parameter space for the precautionary parameter, we find the distributions are often mesokurtic yet bimodal.

With a variance-adjusted estimator, escape dynamics arise instead because the distribution has positive mass at a target inflation rate of zero.

The paper is structured as follows. In Section 2, we review Cho et al. (2002), show that the loss function is asymmetric when expressed as a function of the estimated Phillips curve parameters, and discuss the VALS estimator. In Section 3, we show how the expected loss function varies with the precautionary parameters and explain how unemployment and inflation are impacted by such precautionary learning. We conclude in Section 4.

2. RELATED LITERATURE

As Brainard (1967) highlighted, policymakers face a continual tension between the need to have some confidence in their model in order to make decisions and the reality that, even if they have the correct model, the parameters may change over time. In the environment of Sargent (1999) and Cho et al. (2002), this tension is captured by having the policymaker employ a certainty equivalence rule that ignores any possible parameter uncertainty (either due to nonstationarity or noisy data), reflecting her confidence that the model is right, while also using a constant-gain estimator to update the parameters of the model, which would only be efficient if there is a suspicion that the model is nonstationary, so recent data are more informative than ancient data. This tension has been examined previously and in this section we review that literature and discuss our contribution.

Sims (1988) was among the first to point out this tension, arguing that a central bank ought to see that the data generated by a policy of trying to exploit a Phillips curve with a negative slope will result in a Phillips curve with a vertical slope. Sargent and Williams (2005) sought to reconcile Sims' simulation of a stable Ramsey equilibrium with their previous results in Cho et al. (2002). They concluded that a difference in priors about the model parameters could account for these differing model predictions.

Since frequentist regressions remain the primary workhorse of econometricians, including at central banks, our approach is to see what would happen if the policymaker acknowledges the uncertainty in her parameter estimates and corrects for them in her regression. Tetlow and von zur Muehlen (2004) also study the effects of parameter uncertainty on policy while hewing more closely to Brainard's (1967) framing of the problem. They compare the outcomes of using rules derived from three other methodologies for handling parameter uncertainty: a Bayesian rule, a Knightian uncertainty rule, and a Hansen–Sargent robust control policy. Our VALS estimator is most similar to their Bayesian rule, which also uses covariances to modify inflation targets. Where we differ is that the VALS estimator has a free parameter, the precautionary parameter. When chosen to minimize the econometric loss function, the VALS estimator also reduces average inflation bias, like their Bayesian rule, while staying within the framework of frequentist regressions.

The previous papers and our own all assume that the policymaker considers only one model. Cho and Kasa (2015) consider what happens if the policymaker considers several models simultaneously, including a model with a vertical Phillips curve and no trade-off between inflation and unemployment. Here the policymaker continually runs specification tests to determine if she has been using the best model. In this environment, escapes cause the policymaker to reject every model except the model with the vertical Phillips curve, after which she never switches back to another model. Thus in the long run inflation should converge to the Ramsey rate. Cho and Kasa's specification testing strategy is similar to the recursive approach we propose for modifying the precautionary parameter. In the long run in our environment, the optimal precautionary parameter will tend to explode (either positively or negatively), which an introspective policymaker would likely take as evidence that the model is wrong.

Importantly, our analysis is related to that of McGough (2006), who includes real oil prices in the model of Sargent (1999). Such supply shocks are investigated for their ability to trigger escape-like events. McGough (2006) finds, among other relevant results, that favorable shocks to unemployment decrease the time an economy takes to escape to Nash inflation. Escapes happen much sooner in the presence of such supply shocks.

3. THE MODEL

A central bank chooses its target inflation rate (x_t) to minimize a quadratic loss function defined over realized inflation (π_t) and unemployment (u_t):

$$L = E[\pi_t^2 + \alpha u_t^2], \tag{1}$$

where $\alpha > 0$ is the weight placed on unemployment. The relationship between u_t and π_t is governed by a Phillips curve [i.e., the actual law of motion (ALM)] for the economy:

$$u_t = u_{\text{NR}}^* - \omega^*(\pi_t - \widehat{x}_t) + \sigma_1 W_{1t},\ u_{\text{NR}}^* > 0,\ \omega^* > 0,\ \sigma_1 > 0, \tag{2}$$

in which π_t is determined by

$$\pi_t = x_t + \sigma_2 W_{2t},\ \sigma_2 > 0. \tag{3}$$

The mutually independent shocks $(W_{1t}, W_{2t})'$ are *i.i.d.* and normally distributed with zero means and unit variances. In the Phillips curve (2), $\widehat{x}_t$ is the private sector's expectation of inflation at t and the parameters u_{NR}^* and ω^* are, respectively, the actual natural rate of unemployment and (inverse) slope. Since the public's expectations are rational, $\widehat{x}_t$ will ultimately be set equal to x_t in equilibrium. The outcomes that arise when the central bank knows the parameters u_{NR}^* and ω^* are well known: if x_t is fixed after the public forms $\widehat{x}_t$, the high-inflation Nash equilibrium outcome is a targeted inflation rate of $x_t^* = \alpha\omega^* u_{\text{NR}}^*$; if x_t is fixed before the public forms $\widehat{x}_t$, the low-inflation Ramsey equilibrium outcome would have

https://doi.org/10.1017/S1365100518000731 Published online by Cambridge University Press

the central bank set $x_t^* = 0$. Since the Nash outcome is time-consistent, the central bank suffers from an inflationary bias absent a technology that allows it to credibly commit to the Ramsey rule.

Under adaptive learning, the central bank believes instead in a misspecified version of the Phillips curve:

$$u_t = u_{\mathrm{NR},t} - \omega_t \pi_t + \eta_t, \tag{4}$$

where η_t is a mean-zero noise variable with variance σ_3^2, uncorrelated with the other time-t-dated variables. The intercept and slope terms in this perceived law of motion (PLM) are indexed by t since every period the central bank re-estimates the Phillips curve.

Given (3) and (4), the loss function is

$$L(x_t | u_{\mathrm{NR},t}, \omega_t) = x_t^2 + \alpha(u_{\mathrm{NR},t} - \omega_t x_t)^2 + \alpha\sigma_3^2 + (1 + \alpha\omega_t^2)\sigma_3^2 \tag{5}$$

when expressed as a function of the central bank's policy variable (the inflation target x_t). Given current estimates $\widehat{u}_{\mathrm{NR},t}$ and $\widehat{\omega}_t$, the central bank chooses x_t to minimize $L(x_t | \widehat{u}_{\mathrm{NR},t}, \widehat{\omega}_t)$, which gives the optimal target inflation rate:

$$x(\widehat{\omega}_t, \widehat{u}_{\mathrm{NR},t}) = \frac{\alpha \widehat{\omega}_t \widehat{u}_{\mathrm{NR},t}}{1 + \alpha \widehat{\omega}_t^2}. \tag{6}$$

Note that (6) is nonlinear in $\widehat{\omega}_t$. It is this nonlinearity that drives our results, described in the sections below.

3.1. Precaution in Statistical Decision Making

Following Berger (1985), suppose that a hypothetical statistical decision maker (SDM) is estimating a parameter $\theta \in \mathbf{R}$ so as to minimize the function $f(\theta, \widehat{\theta})$, where $\widehat{\theta}$ is his estimate of θ. A rational SDM operating in a frequentist paradigm will specify an estimator $\widehat{\theta}(y)$, a function of the data y, with the objective of minimizing

$$E\left[f(\theta, \widehat{\theta}(y))\right], \tag{7}$$

where the expectation is taken over y. To ensure that finding the truth is optimal, we assume that f is positive definite so $f(\theta, \widehat{\theta}) \geq 0$ for all $(\theta, \widehat{\theta})$ and $f(\theta, \widehat{\theta}) = 0$ if and only if $\theta = \widehat{\theta}$. Further assuming f is C^4, the positive-definiteness condition implies

$$\frac{\partial f(\widehat{\theta}, \theta)}{\partial \widehat{\theta}} |_{\widehat{\theta} = \theta} = 0 \tag{8}$$

and

$$\frac{\partial^2 f}{\partial \widehat{\theta}^2}(\theta, \theta) \geq 0 \tag{9}$$

for all θ.

https://doi.org/10.1017/S1365100518000731 Published online by Cambridge University Press

Given θ, a Taylor expansion of $f(\theta,\widehat{\theta})$ for $\widehat{\theta}$ near θ is

$$f(\theta,\widehat{\theta}) = f(\theta,\theta) + \frac{1}{2}\frac{\partial^2 f}{\partial\widehat{\theta}^2}(\theta,\theta)(\widehat{\theta}-\theta)^2 + \frac{1}{6}\frac{\partial^3 f}{\partial\widehat{\theta}^3}(\theta,\theta)(\widehat{\theta}-\theta)^3 + O\left((\widehat{\theta}-\theta)^4\right). \tag{10}$$

Thus the SDM wishes to choose a function $\widehat{\theta}(y)$ that minimizes

$$E\left[\frac{1}{2}\frac{\partial^2 f}{\partial\widehat{\theta}^2}(\theta,\theta)(\widehat{\theta}(y)-\theta)^2 + \frac{1}{6}\frac{\partial^3 f}{\partial\widehat{\theta}^3}(\theta,\theta)(\widehat{\theta}(y)-\theta)^3\right],$$

where again the expectation is with respect to y. The resulting first-order condition is

$$\frac{\partial^2 f}{\partial\widehat{\theta}^2}(\theta,\theta)\left(E\left[\widehat{\theta}(y)\right]-\theta\right) + \frac{1}{2}\frac{\partial^3 f}{\partial\widehat{\theta}^3}(\theta,\theta)E\left[(\widehat{\theta}(y)-\theta)^2\right] = 0. \tag{11}$$

If the third derivative of f at $\widehat{\theta}=\theta$ vanishes, (11) reduces to the condition that the SDM should choose a consistent estimator such that $E\left[\widehat{\theta}(y)\right]=\theta$. In the more general case where the third derivative does not vanish, however, an optimal estimator will satisfy

$$E\left[\widehat{\theta}(y)\right] = \theta - \frac{1}{2}\frac{\partial^3 f/\partial\widehat{\theta}^3(\theta,\theta)}{\partial^2 f/\partial\widehat{\theta}^2(\theta,\theta)}E\left[(\widehat{\theta}(y)-\theta)^2\right] \approx \theta - \frac{1}{2}\frac{\partial^3 f/\partial\widehat{\theta}^3(\theta,\theta)}{\partial^2 f/\partial\widehat{\theta}^2(\theta,\theta)}V\left[\widehat{\theta}(y)\right], \tag{12}$$

since, if (12) holds, $E\left[\widehat{\theta}(y)\right]\approx\theta$ for $i=1,\ldots,m$. If $\partial^3 f/\partial\widehat{\theta}^3(\theta,\theta)>0$, the loss function is higher if $\widehat{\theta}(y)=\theta+\varepsilon$ than if $\widehat{\theta}(y)=\theta-\varepsilon$, where $\varepsilon>0$. Consequently, it is optimal for the SDM to bias down his estimate, erring on the side of caution, since it is less costly to underestimate θ than it is to overestimate θ. The more imprecise his estimate $\widehat{\theta}(y)$ is, the larger the bias should be. Conversely, if $\partial^3 f/\partial\widehat{\theta}^3(\theta,\theta)<0$, he should bias his estimates in the opposite direction.

Suppose that $\widehat{\theta}^u(y)$ is an unbiased estimator with variance $\sigma^2(y)$. Then (12) suggests the SDM replace the unbiased estimator with a variance-adjusted estimator

$$\widehat{\theta}_i^v(y) = \widehat{\theta}_i^u(y) + \frac{1}{2}a\sigma^2(y). \tag{13}$$

The parameter a in (13) is what we call the *precautionary parameter*. Optimally, given (12), the SDM should choose

$$a = -\frac{\partial^3 f/\partial\widehat{\theta}^3(\theta,\theta)}{\partial^2 f/\partial\widehat{\theta}^2(\theta,\theta)}. \tag{14}$$

For example, for Varian's (1975) LINEX loss function

$$f(\theta,\widehat{\theta}) = \exp\left(-a(\widehat{\theta}-\theta)\right) + a(\widehat{\theta}-\theta) - 1, \tag{15}$$

the right hand side of (14) is constant and equal to the parameter a.

3.2. Asymmetry in the KPBG Loss Function

What are the properties of the central bank's econometric loss function, assuming it seeks to minimize its loss function (1) when estimating the parameters of the Phillips curve? Since the choice of what policy to set at time t given the latest estimates is not a dynamic problem, we suppress time indices in this section. Let us begin by considering what happens if the policymaker is correct in her beliefs about the economy, so the Phillips curve (4) is the law of motion for the economy instead of (2). In the absence of any misspecifications, the econometric loss function is obtained simply by substituting the optimal inflation target (6), given the current parameter estimates $(\widehat{\omega}, \widehat{u}_{\rm NR})$, into the perceived policy loss function (5), which depends on the true, albeit unknown parameters $(\omega, u_{\rm NR})$:

$$\begin{aligned}\widehat{L}(\widehat{\omega}, \widehat{u}_{\rm NR}, \omega, u_{\rm NR}) &= L\,(x(\widehat{\omega}, \widehat{u}_{\rm NR})|\omega, u_{\rm NR}) \\ &= \left(1+\alpha\omega^2\right)\left(\frac{\alpha\widehat{\omega}\,\widehat{u}_{\rm NR}}{1+\alpha\widehat{\omega}^2}\right)^2 - 2u_{\rm NR}\omega\frac{\alpha^2\widehat{\omega}\,\widehat{u}_{\rm NR}}{1+\alpha\widehat{\omega}^2} \\ &\quad + \alpha(u_{\rm NR}^2+\sigma_1^2) + (1+\alpha\omega^2)\sigma_2^2. \end{aligned} \tag{16}$$

The first partial derivatives of (16) with respect to the estimates are

$$\begin{aligned}\frac{\partial\widehat{L}}{\partial\widehat{\omega}} &= 2(1+\alpha\omega^2)\frac{\alpha^2\widehat{\omega}\,\widehat{u}_{\rm NR}^2}{(1+\alpha\widehat{\omega}^2)^2} - 4(1+\alpha\omega^2)\frac{\alpha^3\widehat{\omega}^3\widehat{u}_{\rm NR}^2}{(1+\alpha\widehat{\omega}^2)^3} \\ &\quad - 2u_{\rm NR}\omega\frac{\alpha^2\widehat{u}_{\rm NR}}{1+\alpha\widehat{\omega}^2} + 4u_{\rm NR}\omega\frac{\alpha^3\widehat{\omega}^2\,\widehat{u}_{\rm NR}}{(1+\alpha\widehat{\omega}^2)^2}, \\ \frac{\partial\widehat{L}}{\partial\widehat{u}_{\rm NR}} &= 2\frac{1+\alpha\omega^2}{(1+\alpha\widehat{\omega}^2)^2}\alpha^2\widehat{\omega}^2\widehat{u}_{\rm NR} - 2u_{\rm NR}\frac{\alpha^2\widehat{\omega}\omega}{1+\alpha\widehat{\omega}^2}. \end{aligned} \tag{17}$$

If $\widehat{\omega} = \omega$ and $\widehat{u}_{\rm NR} = u_{\rm NR}$,

$$\frac{\partial\widehat{L}}{\partial\widehat{\omega}}(\omega, u_{\rm NR}) = 2\frac{\alpha^2\omega u_{\rm NR}^2}{1+\alpha\omega^2} - 2\frac{\alpha^2\omega u_{\rm NR}^2}{1+\alpha\omega^2} - \frac{4\alpha^3\omega^3u_{\rm NR}^2}{(1+\alpha\omega)^2} + \frac{4\alpha^3\omega^3u_{\rm NR}^2}{(1+\alpha\omega^2)^2} = 0,$$

$$\frac{\partial\widehat{L}}{\partial\widehat{u}_{\rm NR}}(\omega, u_{\rm NR}) = \frac{2\alpha^2\omega^2u_{\rm NR}}{1+\alpha\omega^2} - 2u_{\rm NR}\frac{\alpha^2\omega^2}{1+\alpha\omega^2} = 0.$$

Since the loss function is quadratic in $\widehat{u}_{\rm NR}$,

$$\frac{\partial^2\widehat{L}}{\partial\widehat{u}_{\rm NR}^2}(\omega, u_{\rm NR}) = \frac{2\alpha^2\omega^2}{1+\alpha\omega^2} \geq 0,$$

and all higher-order derivatives with respect to $\widehat{u}_{\rm NR}$ vanish, it is optimal that the central bank should learn the truth about $u_{\rm NR}$, and a consistent estimator of $u_{\rm NR}$ is optimal. Thus in the following we can set $\widehat{u}_{\rm NR} = u_{\rm NR}$. This simplifies the partial derivative with respect to $\widehat{\omega}$, (17), which is then exactly proportional to $\alpha^2 u_{\rm NR}^2$:

https://doi.org/10.1017/S1365100518000731 Published online by Cambridge University Press

$$\frac{1}{\alpha^2 u_{\rm NR}^2}\frac{\partial\widehat{L}}{\partial\widehat{\omega}} = 2(1+\alpha\omega^2)\frac{\widehat{\omega}}{(1+\alpha\widehat{\omega}^2)^2} - 4(1+\alpha\omega^2)\frac{\alpha\widehat{\omega}^3}{(1+\alpha\widehat{\omega}^2)^3} - \frac{2\omega}{1+\alpha\widehat{\omega}^2} + \frac{4\alpha\omega\widehat{\omega}^2}{(1+\alpha\widehat{\omega}^2)^2}.$$

The second derivative of the loss function is then given by

$$\frac{1}{\alpha^2 u_{\rm NR}^2}\frac{\partial^2\widehat{L}}{\partial\widehat{\omega}^2}(\omega) = \frac{2+4\alpha\omega^2+2\alpha^2\omega^4-8\alpha\omega^2}{(1+\alpha\omega^2)^3} = \frac{2-4\alpha\omega^2+2\alpha^2\omega^4}{(1+\alpha\omega^2)^3} = \frac{2(1-\alpha\omega^2)^2}{(1+\alpha\omega^2)^3} \geq 0,$$

which is strictly positive as long as $\omega \neq \alpha^{-1/2}$. The third derivative is

$$\frac{1}{\alpha^2 u_{\rm NR}^2}\frac{\partial^3\widehat{L}}{\partial\widehat{\omega}^3}(\omega) = -\frac{12\alpha\omega}{(1+\alpha\omega^2)^4}(3-\alpha\omega^2)(1-\alpha\omega^2), \quad \textbf{(18)}$$

which only vanishes for the knife-edge cases where $\alpha\omega^2 = 1$ or $\alpha\omega^2 = 3$. This means there generally will be an asymmetry in the econometric loss function for the Phillips curve slope, and optimal estimation will require a biased estimator for ω. The sign of the asymmetry is ambiguous and depends on the true Phillips curve slope. Because of this potential for asymmetry, the optimal precautionary parameter a, given by (14) in the previous section, is

$$a^* = -\frac{\frac{\partial^3\widehat{L}}{\partial\widehat{\omega}^3}}{\frac{\partial^2\widehat{L}}{\partial\widehat{\omega}^2}} = -\frac{-\frac{12\alpha\omega}{(1+\alpha\omega^2)^4}(3-\alpha\omega^2)(1-\alpha\omega^2)}{\frac{2(1-\alpha\omega^2)^2}{(1+\alpha\omega^2)^3}} = \frac{6\alpha\omega(3-\alpha\omega^2)}{1-\alpha^2\omega^2}. \quad \textbf{(19)}$$

Thus, the policymaker should only believe that least-squares learning is efficient for the special case where $\alpha\omega^2 = 3$. Note also that the optimal bias could potentially be unbounded if $\alpha^2\omega^2$ approaches 1.

3.3. The VALS Estimator

Under the constant-gain ROLS learning algorithm, at each period t the estimates $\xi_t = (\widehat{u}_{{\rm NR}_t}, \widehat{\omega}_t)'$ are updated using the least-squares formula:

$$\xi_t = (Z_t' Z_t)^{-1} Z_t' U_t, \quad \textbf{(20)}$$

where $U_t = (u_0, \ldots, u_{t-1})'$ and $Z_t = [\mathbf{1}'\,(\pi_0, \ldots, \pi_{t-1})']$. Let $R_t = g Z_t' Z_t$ be the matrix of second moments where g is the gain. Defining $z_t = (1, \pi_t)'$, the estimates and R_t evolve according to

$$\xi_t = \xi_{t-1} + g R_t^{-1} z_{t-1}(u_{t-1} - z_{t-1}'\xi_{t-1}), \quad \textbf{(21)}$$

$$R_t = R_{t-1} + g\left(z_{t-1}z_{t-1}' - R_{t-1}\right), \quad \textbf{(22)}$$

in which, unbeknownst to the central bank, the u_{t-1} in the evolution of ξ_t above is determined by the ALM. The gain (g) can be interpreted as the inverse of the

https://doi.org/10.1017/S1365100518000731 Published online by Cambridge University Press

learning horizon T.[7] A constant-gain learning algorithm is appropriate if the policymaker is concerned about structural change to the Phillips curve and believes data from say, 100 years ago, is not informative about what is happening today.

The constant-gain ROLS estimator departs from a decreasing gain algorithm in that escape dynamics can arise more frequently, as discussed in Section 1. That is, if by chance the Phillips curve in the last T periods appears flat, the inflation rate can temporarily drop to the low value of the Ramsey equilibrium. It is this constant-gain ROLS learning case that we modify here. The first step in our analysis is the assumption of "precaution" in the interpretation of data and so the use of a constant-gain VALS estimator of $(u_{\mathrm{NR_t}}, \omega_t)'$. This estimator, denoted by ξ_t^a, is defined by

$$\xi_t^a = \xi_t + \frac{1}{2}\widehat{\Sigma}_t a, \tag{23}$$

where $\widehat{\Sigma}_t$ is a constant-gain estimator of the variance–covariance matrix for the least-squares estimator ξ_t. The vector $a = (a_{u_{\mathrm{NR}}}, a_\omega)'$ collects the precautionary *parameters* that determine whether the estimator biases up or down the two Phillips curve parameters u_{NR} and ω. If either $a_{u_{\mathrm{NR}}}$ or a_ω is negative (positive), then the bias for that variable is against overestimation (underestimation).

The variance-covariance matrix is defined by

$$\Sigma_t = t^{-1} R_t^{-1} \sigma_t^2, \tag{24}$$

where

$$\sigma_t^2 = \frac{1}{t-1}\sum_{i=0}^{t-1}(u_i - z_i'\xi_t)^2 \tag{25}$$

is the most recent estimate of the variance of η_t, the noise in the PLM. Expanding the square in (25), we get

$$\sigma_t^2 = \frac{1}{t-1}\sum_{i=0}^{t-1}\Big[(u_i - z_i'\xi_{t-1})^2 + 2(u_i - \xi_{t-1}' z_i) z_i'(\xi_{t-1} - \xi_t) + (\xi_{t-1} - \xi_t)' z_i z_i'(\xi_{t-1} - \xi_t)\Big]. \tag{26}$$

We have already introduced R_t, which is a constant-gain estimator of $\frac{1}{t}\sum_{i=0}^{t-1} z_i z_i'$. Similarly, we now introduce S_t as a constant-gain estimator of

$$\frac{1}{t}\sum_{i=0}^{t-1} u_i z_i'.$$

We define S_t by the difference equation:

$$S_t = S_{t-1} + g(u_{t-1}x_{t-1}' - S_{t-1}). \tag{27}$$

Then a constant-gain estimator of σ_t^2 can be defined by the difference equation:

$$\widehat{\sigma}_t^2 = \widehat{\sigma}_{t-1}^2 + \frac{g}{1-g}\Big[(u_{t-1} - z'_{t-1}\xi_{t-1})^2 \left(1 + g z'_{t-1} R_t^{-1} z_{t+1}\right) + 2\left(\xi'_{t-1} R_t - S_t\right) R_t^{-1} z_{t-1}(u_{t-1} - z'_{t-1}\xi_{t-1})\Big], \tag{28}$$

and

$$\widehat{\Sigma}_t = g R_t^{-1} \widehat{\sigma}_t^2. \tag{29}$$

4. SIMULATION RESULTS

4.1. Loss Function Surfaces

Given an environment in which agents adaptively learn, the usual procedure is to stochastically approximate the system of equations that characterize model dynamics, including the Ricatti equations associated with ROLS regressions. Such an approximation for "mean dynamics" delivers a system of differential equations that can be examined for stability of model equilibria. Cho et al. (2002) defined escape dynamics as well for the case of ROLS learning under constant gain. Since this analytical procedure is not easily generalized to account for the dynamics of the variance, we instead opt to provide evidence via simulation on whether Ramsey or Nash inflation is obtained under precautionary learning. The simulation results prompt the intuition described above: precautionary learning, instantiated by the use of a VALS estimator on the part of a central bank, can lead to a higher frequency of escape and a lowering of the long-run average rate of inflation.

The key to this intuition is the nature of the loss function (L) surface as we vary the precautionary parameter values, discussed below. We note that we are plotting the loss function as a function of the precautionary parameters that inflation and unemployment implicitly depend upon; the expected loss function is computed as the average loss over a simulation of 10,000 periods.

We generate each of these figures (and all reported simulations below) having fixed α at 1, ω at 2, and u_{NR} at 5, so that the Nash level of inflation from $x_t^* = \alpha\omega^* u_{\mathrm{NR}}^*$ is 10 while the Ramsey equilibrium inflation rate is zero. The decreasing gain learning algorithm will converge to the Nash equilibrium outcome in which the Phillips curve is perceived to have the actual slope of 2 and an intercept of 25, which adds to u_{NR} the effect of the public's expectations about inflation. The standard deviations of the shocks, σ_1 and σ_2, are set to 1.5. We start each simulation looking at data generated by a predecessor policymaker who targeted the inflation rate at its Nash value.

In our baseline calibration, we set the horizon for the constant-gain algorithm to $T = 20$ periods, implying a gain of $g = \frac{1}{T} = 0.05$. The loss surface is shown in Figure 1. We also use this calibration for the time series plots of targeted inflation that we report below.

https://doi.org/10.1017/S1365100518000731 Published online by Cambridge University Press

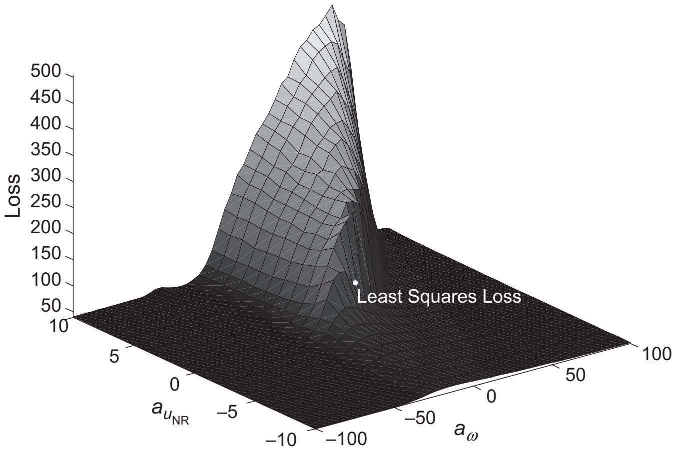

FIGURE 1. Loss function [$L(x_t)$, 0.05 gain].

A high gain translates into a shorter time span of data employed in learning the Phillips curve parameters. Thus the probability of observing a history in which the Phillips curve is flat is high, raising the probability of escape. Given our interest in examining dynamics with constant-gain VALS versus ROLS learning, a high gain of 0.05 is a natural choice.

In Figure 1, we see that least-squares learning is close to a ridge line of the loss function.[8] Allowing for a bias in the estimation of either of the two Phillips curve parameters [i.e., nonzero values for $(a_{u_{\mathrm{NR}}}, a_\omega)'$] suggests that better policy may result from overestimating the slope and underestimating the intercept relative to least squares. However, the improvement that comes from moving away from least squares in this direction is a purely local phenomenon. Whereas slight underestimates of the slope lead to a worsening of policy, large underestimates improve policy, and likewise for overestimates of the intercept. The policymaker will be nearly indifferent between receiving very high or very low estimates of the parameters as long as she does not receive estimates in the vicinity of least squares.

4.2. Time Series Simulations

Why do biased estimates of the Phillips curve parameters improve policy? Do they cause the policymaker to target rates of inflation lower than those consistent with a Nash outcome? We investigate these questions by looking at the time series generated by our model as we vary the values for the precautionary parameters, holding the shock processes and the value of the gain constant. For our baseline parameters, the period loss function for the Ramsey equilibrium is 38. In contrast, the loss function for the Nash equilibrium is the much higher value of 138. First, we consider what happens for the familiar case of constant-gain ROLS learning, in

https://doi.org/10.1017/S1365100518000731 Published online by Cambridge University Press

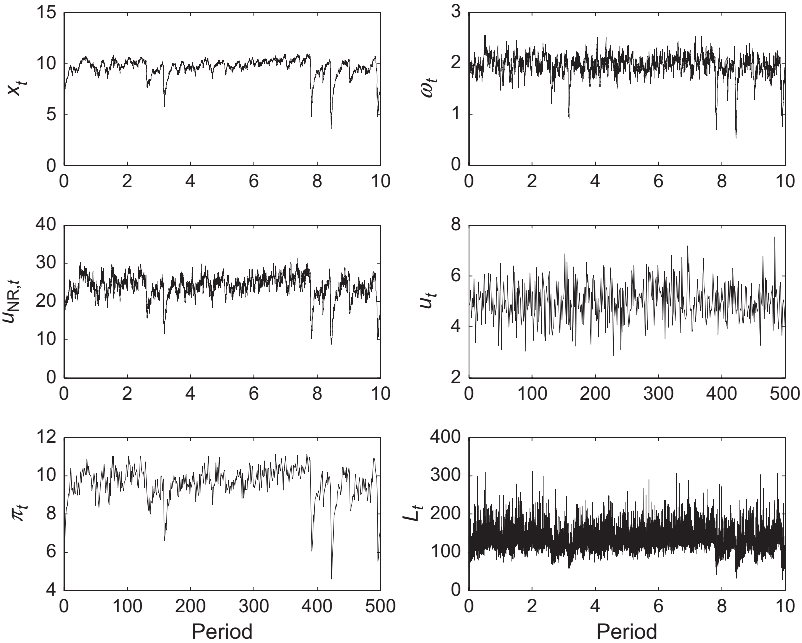

FIGURE 2. Time paths under constant-gain ROLS learning.

which $(a_\omega, a_{u_{NR}}) = (0, 0)$. The time series for targeted inflation and the estimated Phillips curve parameters are shown in Figure 2.[9]

Consistent with Cho et al. (2002), the inflation target, shown in the top panel, remains around the Nash equilibrium value of 10 but there are "escapes" where the target falls close to the Ramsey level of zero. The plots of ω_t and u_{NR_t} in the second and third panels reflect the intuition of escape, as follows. The average estimates of ω and u_{NR} are 1.9520 and 23.8183, respectively, both close to the values that would be obtained by regressing u_t on π_t according to (4) in the Nash equilibrium.[10] However, a string of data occurs in which the correlation of u_t and π_t falls, resulting in smaller estimates of both ω and u_{NR}. The policymaker thinks there is a smaller trade-off between unemployment and inflation, so according to the targeted rate under the PLM

$$x_t^{PLM} = \frac{\alpha \omega_t u_{NR,t}}{1 + \alpha \omega_t^2}, \tag{30}$$

she reduces the inflation target. As a consequence of the lower target, the slope and intercept estimates remain depressed, so it takes time for the system to return to Nash equilibrium values. The period loss function falls during the escape, but most of the time it is around the Nash equilibrium value of 138.

We note that in Figures 2 and 3, while we observe sharp and deep drops indicative of escapes, the analysis of Kolyuzhnov et al. (2014) may apply. That is,

https://doi.org/10.1017/S1365100518000731 Published online by Cambridge University Press

usually the gain has to be very small for standard large deviation approximations to accurately characterize escapes.[11] The gain we employ is larger and consistent with Cho et al. (2002) but the difference in parameter values reinforces the issues raised in Kolyuzhnov et al. (2014), namely that the gap between Nash and Ramsey inflation is widened. However, since Evans and Honkapohja (2001, p. 354) we know that the mean dynamics of such systems point to Ramsey for only small perturbations around Nash. In our analysis what is important is the drop in targeted inflation in our VALS learning, a feature we focus on in our simulation results.

Next we consider what happens as we vary the precautionary parameters on a grid away from $(0, 0)$. We use the same sequence of shocks as those that generated the time series reported in Figure 2. Under rational expectations the ALM (2) eliminates any dependence of unemployment on the central bank's policy decisions, so the unemployment series is purely a function of the shocks. Consequently, all of the following plots have the same unemployment series as in Figure 2. Only the series for the target inflation rate and the realized inflation rate will depend on the central bank's policy decisions, which themselves depend on the precautionary parameters. Since the realized inflation rate differs from the target inflation only by the addition of noise, we only plot targeted inflation.[12]

Intuitively, one might expect the policymaker to achieve better results if the slope is underestimated since that will increase the probability of finding there is no trade-off between unemployment and inflation, producing a Ramsey outcome. The time series for targeted inflation that result from setting $a_\omega < 0$ are shown in Figure 3.[13] However, whether policy will improve if a_ω is increased above or below zero depends on the parameters of the model. Differentiating (30) by ω, we find

$$\frac{\partial x_t^{\text{PLM}}}{\partial \omega_t} = \frac{\alpha u_{\text{NR},t}(1 - \alpha\omega_t^2)}{(1 + \alpha\omega_t^2)^2}, \tag{31}$$

which is of ambiguous sign. For our values, $\alpha\omega^2 = 4 > 1$, where ω is the actual slope of the Phillips curve. Thus a small decrease in the slope estimate will actually increase the target inflation rate.

Consistent with our expectation, escapes occur more frequently as a_ω drops below zero in Figure 3. However, because the inflation rate is higher when the economy is not escaping, the loss function increases as we decrease a_ω for small $|a_\omega|$. When $a_\omega = -10$, the economy escapes away from the Nash outcome about half the time. The period loss function can fall almost to zero during these many escapes, but these gains are dominated by the higher values of the period loss function that arise during normal periods when inflation is targeted around 15 instead of 10. Overall, the loss with $a_\omega = -10$ averages to 152, which is higher than the loss of 138 if the economy stays at the Nash equilibrium. On the other hand, when $\alpha_\omega = -20$, escapes become so frequent that the economy never has a chance to converge to a steady state. These frequent escapes yield a loss of 80 that is much less than the loss under the Nash equilibrium. In pushing a_ω so low,

TABLE 1. Simulation statistics ($a_\omega \lesseqgtr 0$, $a_{u_{\rm NR}} = 0$)

	x_t		ω_t		$u_{{\rm NR},t}$		Loss	
a_ω	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
0	9.5941	0.8625	1.9520	0.2210	23.8183	3.0786	131.5671	28.5367
−10	9.7048	4.4372	1.2522	0.4992	21.6188	9.2691	152.7371	78.5269
−20	5.0913	4.4372	0.5841	0.4145	12.3178	6.3095	80.6732	56.1517
−100	−1.1090	0.4326	−1.0949	0.8120	2.9258	0.4085	40.1454	41.3213
10	7.0840	0.4692	2.2232	0.1949	18.9694	1.7265	89.1978	26.5272
20	5.4414	0.5109	2.4699	0.2395	15.6257	1.4618	68.6634	28.3052
100	1.8490	0.3866	4.5875	0.8745	8.6088	0.7137	42.3360	34.4630

Note: The table documents the mean and standard deviation of each of the series (x_t, ω_t, $u_{{\rm NR},t}$, and the Loss) as a_ω is varied between 0 and 100, while keeping $a_{u_{\rm NR}}$ fixed at 0.

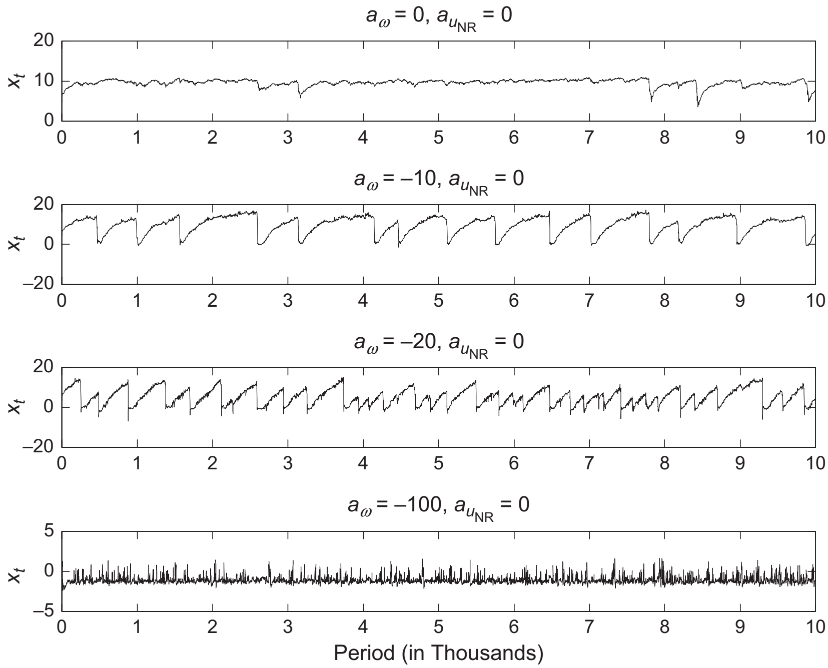

FIGURE 3. Targeted inflation ($a_\omega < 0$, $a_{u_{\rm NR}} = 0$).

we have passed over the ridge line of the loss surface in Figure 1. Decreasing a_ω even further can push the loss down almost to its Ramsey equilibrium value.[14] In the bottom panel of Figure 3, we see that for $a_\omega = -100$, target inflation is quite steady at a value slightly less than 0. Summary statistics for target inflation, actual inflation, unemployment, and the period loss function are summarized in Table 1.

https://doi.org/10.1017/S1365100518000731 Published online by Cambridge University Press

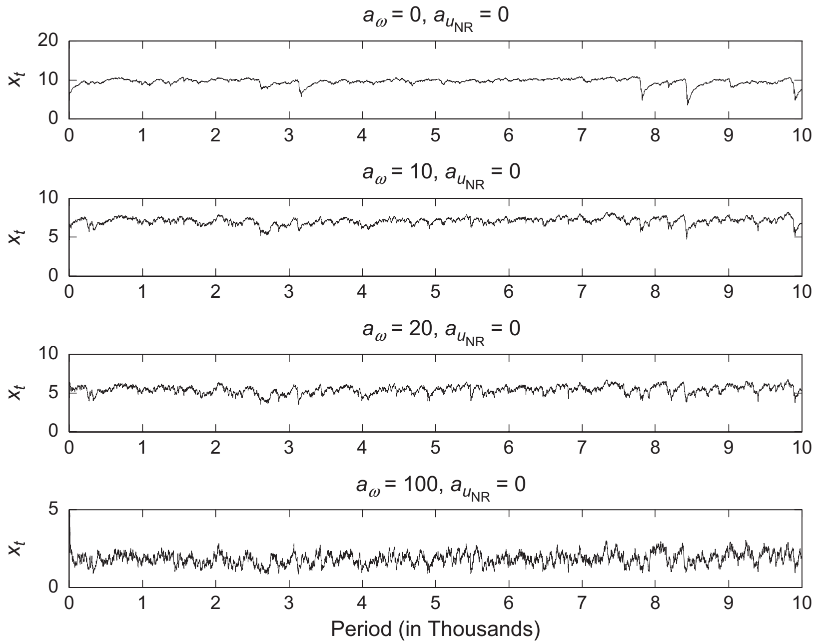

FIGURE 4. Targeted inflation ($a_\omega > 0$, $a_{u_{NR}} = 0$).

In contrast, when the policymaker overestimates the Phillips curve slope as in Figure 4, she thinks the trade-off between unemployment and inflation is so large that she does not need much inflation to achieve the desired low unemployment. Escapes never happen in these simulations because she never convinces herself that the trade-off has gone away. Instead, target inflation remains fairly constant at a value that decreases with a_ω. Consequently the loss also decreases with a_ω for $a_\omega \geq 0$ (over the interval $a_\omega \in [-100, 100]$ that we investigated). When $a_\omega = 20$, target inflation averages to 5.4, and the period loss function averages to 68. As a_ω increases all the way to 100 the average target inflation falls to 1.85, and the period loss function averages to 42, which is nearly the value of the Ramsey equilibrium. Indeed, the economy behaves quite similarly for both $a_\omega = \pm 100$.

A precautionary bias can also be applied to estimation of the intercept of the PLM, that is, the natural rate of unemployment. What happens then? What of VALS learning on the intercept of the PLM? Overestimation of the intercept leads indirectly to lower values of the slope estimate. As a consequence, the probability of escape increases. However, if we differentiate (30) by u_{NR}, we find

$$\frac{\partial x_t^{PLM}}{\partial u_{NR,t}} = \frac{\alpha \omega_t}{1 + \alpha \omega_t^2} > 0, \tag{32}$$

so overestimation of the intercept also produces higher values for target inflation.[15] The benefit of more frequent escapes will be overwhelmed by the cost

https://doi.org/10.1017/S1365100518000731 Published online by Cambridge University Press

TABLE 2. Simulation statistics ($a_\omega = 0$, $a_{u_{NR}} \lesseqgtr 0$)

	x_t		ω_t		$u_{NR,t}$		Loss	
$a_{u_{NR}}$	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
0	9.5941	0.8625	1.9520	0.2210	23.8183	3.0786	131.5671	28.5367
−1	7.0459	0.6250	1.9422	0.2182	17.4322	2.1946	88.8257	27.0004
−10	3.1579	0.6051	1.8618	0.2371	7.6556	1.6542	49.1406	31.9521
1	11.0747	6.9616	1.1186	0.4878	24.0052	13.5953	209.9474	175.6188
10	9.8065	8.1298	0.3390	0.1963	32.9897	26.9330	200.9538	217.6791
min L	−0.7486	0.5693	−0.6599	0.6202	2.6151	0.6063	39.3381	40.3295

Notes: The table documents the mean and standard deviation of each of the series (x_t, ω_t, $u_{NR,t}$, and the Loss) as $a_{u_{NR}}$ is varied between 0 and 100, while keeping a_ω fixed at 0. The last row of the table provides these statistics for $(a_{U_{NR}}, a_\omega) = (-10, -76)$.

of higher inflation during normal times for small $a_{U_{NR}}$, and the period loss function averages to 210 for $a_{U_{NR}} = 1$. In contrast, Figure 4(b) in the Supplementary Material shows that escapes are eliminated when we underestimate the intercept, but target inflation averages to a lower value of 7.04 when $a_{U_{NR}} = -1$. In that case, the period loss function averages to 88. Thus variation around $a_{u_{NR}} = 0$ suggests underestimation of the intercept. This pattern persists if we increase $|a_{U_{NR}}|$ even further as documented in Table 2.

Finally, Figure 5 shows the time series with the parameters $(a_{U_{NR}}, a_\omega) = (-10, -76)$ that produce the minimum loss function value of 39.3381 for the region of the parameter space over which we conducted the simulations. The last row of Table 2 provides the associated statistics for the variables of interest in this case. For this last specification of $a_{U_{NR}}$ and a_ω, which produces the minimum loss, we are across the ridge line from ROLS learning, and the targeted rate of inflation averages to −0.7486 (a slight deflation). This suggests that a monetary policymaker who learns about the Phillips curve using a constant-gain VALS estimator can attain a Ramsey-like equilibrium, in contrast to what happens if she restricts herself to an ROLS estimator as in Cho et al. (2002) in which the precautionary parameters are implicitly set to zero.

4.3. Escape Probabilities

The analysis above provides simulation results of the model dynamics as the value of precautionary parameters are altered while holding other parameters fixed. However, some additional considerations relate to escape probabilities in the model.[16] We calculated numerically the stationary distributions of estimators, and the target inflation, in order to understand escape probabilities under constant gains. To begin this discussion, Figure 6 shows the distribution of the target inflation under ROLS learning with no precaution. The distribution is centered at the Nash value of targeted inflation. Escapes occur because the lower tail of the

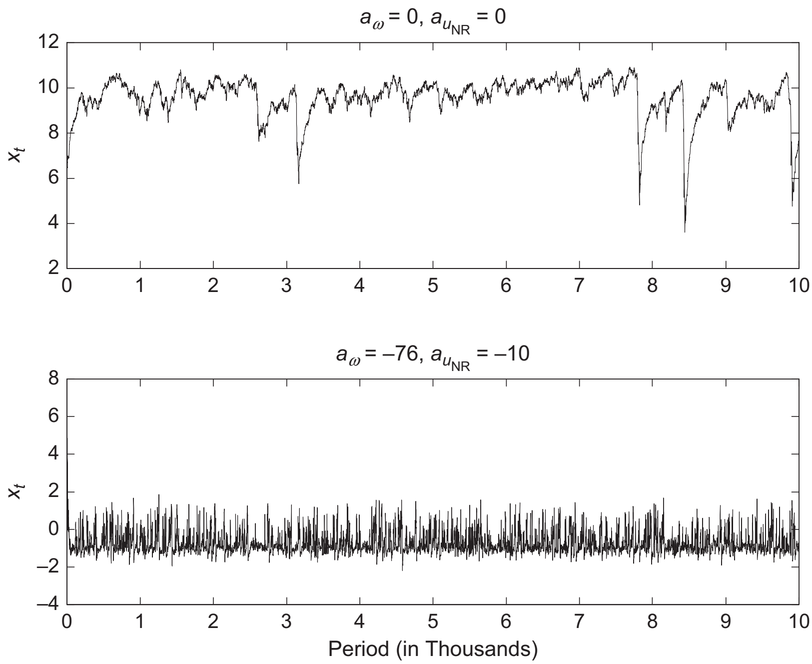

FIGURE 5. Targeted inflation ($a_{u_{NR}} = -10$, $a_\omega = -76$).

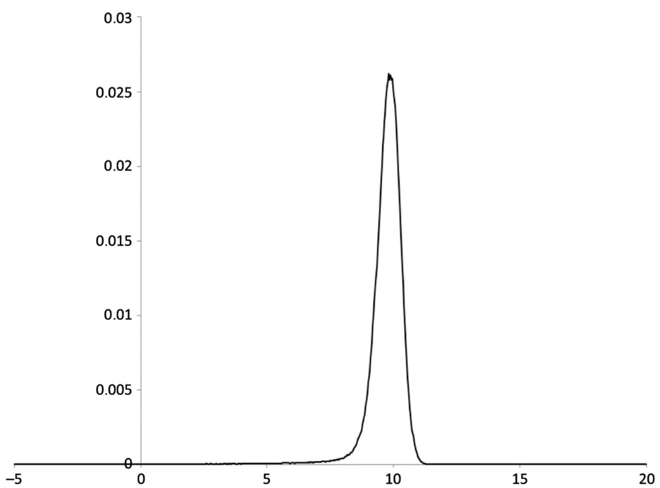

FIGURE 6. Stationary distribution of targeted inflation, no precaution.

https://doi.org/10.1017/S1365100518000731 Published online by Cambridge University Press

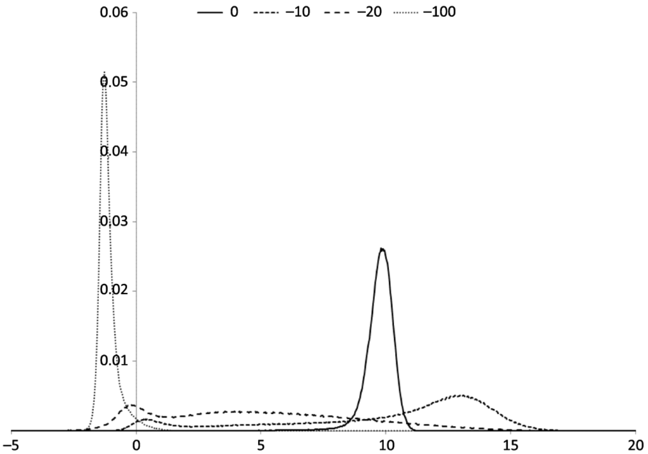

FIGURE 7. Stationary distribution of targeted inflation as slope precaution varied I.

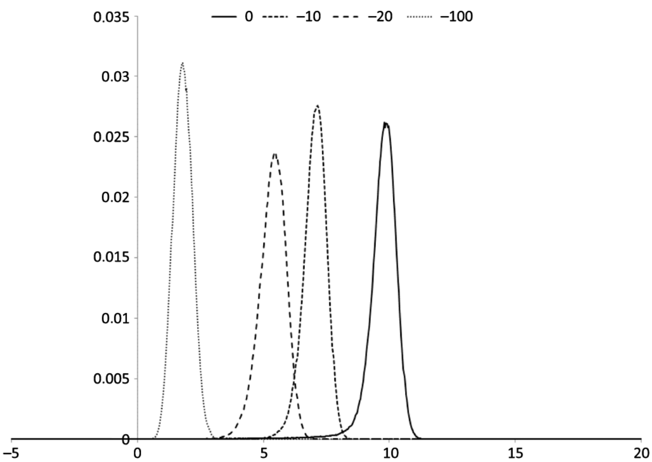

FIGURE 8. Stationary distribution of targeted inflation as slope precaution varied II.

distribution is thick. There is a nonnegligible probability that target inflation will be much lower than the Nash value, in the vicinity of the Ramsey equilibrium that targets zero inflation.

Figures 7–11 demonstrate what happens to the distribution of target inflation as we vary the precautionary parameters. Our first set of plots describe the distribution of targeted inflation as a_ω is varied in the negative direction while keeping the precautionary parameter associated with the intercept in the PLM at zero. We have seen before that as $-a_\omega$ increases in the negative direction, the probability of

https://doi.org/10.1017/S1365100518000731 Published online by Cambridge University Press

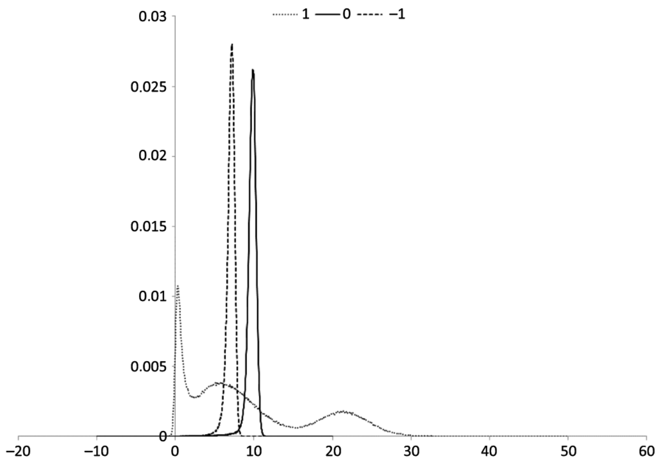

FIGURE 9. Stationary distribution of targeted inflation as intercept precaution varied.

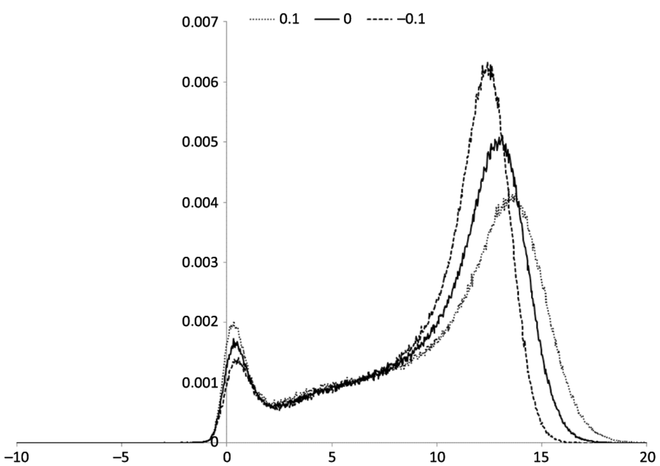

FIGURE 10. Stationary distribution of targeted inflation, nonzero slope precaution, intercept precaution varied I.

escape increases. Naively, one might expect this would show up in the distribution of targeted inflation as a thicker tail, but that is not what happens. Instead the stationary distribution of targeted inflation comes to center at values lower than Ramsey inflation, and a secondary mass develops near zero. As we increase $|a_\omega|$ to 100, the secondary mass eventually predominates. The second plot, Figure 8, shows what happens as a_ω decreases above zero. In this direction, the distribution becomes centered at lower values of targeted inflation while the lower tail gets thinner.

https://doi.org/10.1017/S1365100518000731 Published online by Cambridge University Press

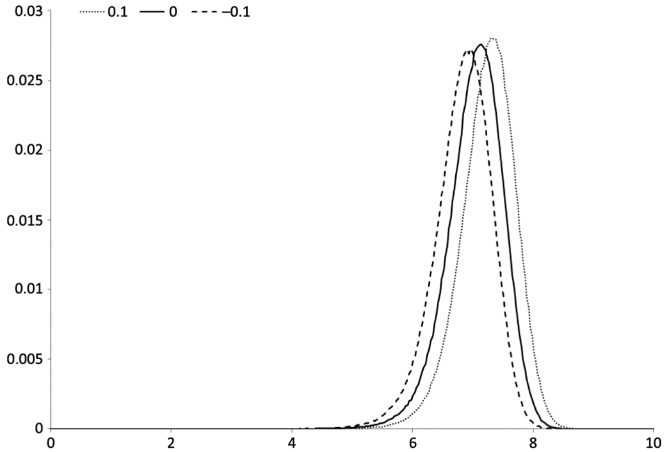

FIGURE 11. Stationary distribution of targeted inflation, nonzero slope precaution, intercept precaution varied II.

Next, we vary only the precautionary parameter associated with the intercept and plot the resulting distributions of targeted inflation in Figure 8. We find that deviation from the least-squares value of 0 does generate movement away from Nash inflation.

Finally, we set the precaution parameter on the slope of the Philips curve to nonzero values and vary the same for the intercept. In Figure 9, we set a_{ω} to 10, vary $a_{U_{\mathrm{NR}}}$, and plot the resulting distributions of targeted inflation. In Figure 10, we set a_{ω} to -10, vary $a_{U_{\mathrm{NR}}}$, and plot the resulting distributions of targeted inflation. In the former case, we find some mass on Ramsey inflation and in the latter case we find movement of the distribution away from Nash inflation.

4.4. On the Value of the Precautionary Parameters

What should the values of the precautionary parameters take in reality? This topic is similar to one which attempts to discover what the value of a constant gain should be in adaptive learning frameworks that analyze, for instance, asset pricing dynamics using S&P 500 price and dividend data.[17] While this is an empirical issue beyond the scope of this paper, we note that prior to even implementing empirics, whether and how a particular learning algorithm is to be motivated in self-referential macroeconomic systems is of first-order importance [see Evans et al. (2010) who suggest that many algorithms are possible]. We motivated the use of a VALS algorithm given the potential asymmetry in an objective function of the policymaker in the otherwise standard endogenous tracking framework of Cho et al. (2002). That is, the asymmetry suggests that higher-order derivatives matter, prompting the use of a VALS algorithm. We presented simulation evidence above that suggests that the use of such an algorithm can reduce long-run inflation and that a policymaker's precaution in learning about fundamentals could lead to higher welfare. These results were generated by varying the values of the precautionary *parameters* over a range. However, another investigation is possible by varying the values of the precautionary parameters over time.

We now let the vector a be a function of time (vs. a constant in the previous sub-sections). The technical aspects of the evolution of the vector a_t are specified in the algorithm provided in the Supplementary Material and we describe it as follows. Recall the optimal a given by (19) and realize that the policymaker believes that she has a correctly specified model, but she also knows from Section 3.2 that the cost of overestimating the slope is different from the cost of underestimating the slope. Policy should dictate that she should include a precautionary bias when estimating the slope, but there is uncertainty about the optimal bias. Inserting the value of the slope ω, which is 2 in the baseline calibration, into (19) suggests that if our model is properly specified, the bias should be about 0.8.[18] We conduct the following "experiment." Employ a bias of -0.8 for 10 years, and then try a bias of 0.8 for another 10 years, and compare the resulting loss to what obtains with no bias and ordinary least-squares learning. If there is statistically significant evidence that the loss is lower with no bias, stay with ordinary least-squares learning indefinitely. If these first experiments are inconclusive, repeat the trial. If, however, there is statistically significant evidence that the loss is lower with a bias, then further such experiments are needed to determine the nonzero bias to employ.

Figure 12 plots the loss and the precautionary parameter on the slope as they evolve over time in the algorithm outlined above with an experimental period of $T = 120$ months and $\phi^* = 1.64$ (with ϕ^* being the relevant critical value). The plots include the error bands and clearly indicate a movement away from a least-squares value for a_ω. This exercise demonstrates that it is entirely possible that the policymaker will find it useful to employ a VALS estimator since the evolution of a_t is away from the ordinary least-squares value of zero (with attendant lower loss). After 10,000 periods, the policymaker would revert to ROLS in only 16% of simulations.

We therefore have provided two sets of results. The first based on varying the values of the precautionary parameters in the calibration. The second by letting the parameters be a function of time and setting out an algorithm for their evolution based on tests that a hypothetical policymaker could conduct. Both sets of results point out that the consequences of precautionary learning are to lower losses and average values of inflation. The latter exercise can be construed to be a stretch on the information that a policymaker may choose to employ or not and does raise a number of questions. However, our motivation in conducting both exercises is to highlight another contribution: our "precautionary" learning results mimic those of Brainard (1967)'s Conservatism Principle in which a policymaker acting in a Bayesian fashion would treat parameter estimates as imprecisely estimated. However, the key difference in our analysis is that the adaptive learning environment is frequentist in nature, that it delivers results similar to those of Brainard (1967) is something we view as our contribution: even in frequentist

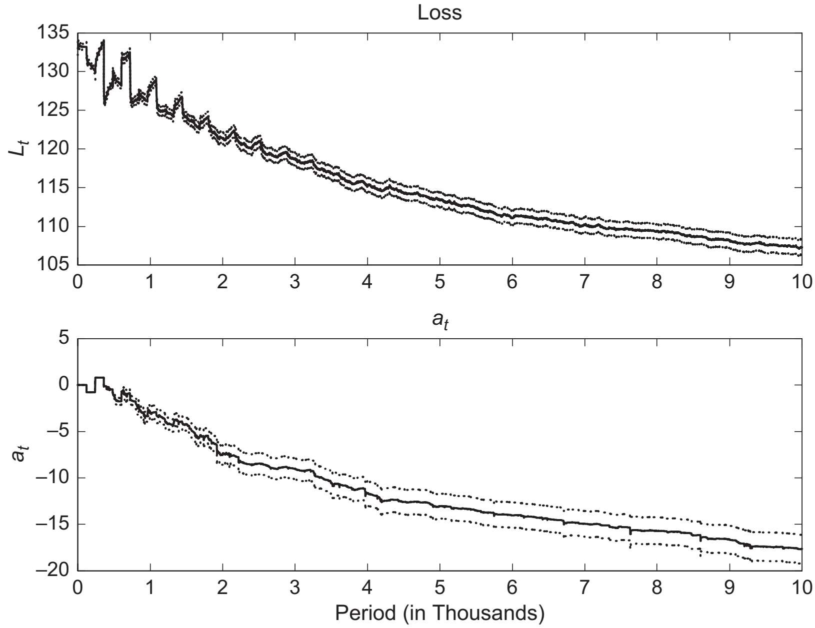

FIGURE 12. Time varying precautionary parameter for slope.

environments a policymaker may internalize parameter uncertainty via use of a variance adjusted learning algorithm and thereby act in a manner to that predicted by Brainard (1967). Both sets of our simulations support this view.

5. CONCLUSION

In replacing expectations with regressions in canonical monetary policy models, escape from a high inflation steady state is possible if a constant-gain recursive least-squares estimator is employed by a monetary authority. In this paper we note that given the reaction function of a policymaker the loss function is asymmetric, so overestimates and underestimates are not equally costly. We then explore the escape dynamics that result if the policymaker uses a VALS that biases her estimates in the more advantageous direction. Our results point toward situations in which such "precaution" on the part of the policymaker could result in the lowering of the long-run rate of inflation in an economy characterized by the time inconsistency of discretionary policy.

Specifically, given the analytically intractable nature of the model, we focus on simulations that indicate that if there is overestimation of the slope relative to least squares, we see an increased frequency of the escapes analyzed by Cho et al. (2002). This is because estimates of the slope of the PLM (a *misspecified* Phillips curve) fall often, along with wide variation in estimates of the intercept. This leads to overshooting dynamics in the targeted rate of inflation. However, this

https://doi.org/10.1017/S1365100518000731 Published online by Cambridge University Press

behavior is not monotonic. For high degrees of slope underestimation, negative slope estimates push down estimates of the natural rate to near zero leading to low values of the targeted rate of inflation. These complex learning dynamics are rich enough to warrant interpretation since they suggest a possible tension between the pull of time consistency and the degree to which a policymaker is cautious about interpreting information, reflected in the estimator employed. Our favored interpretation is that the dynamics can reflect the Conservatism Principle of Brainard (1967) and our contribution is in part to show exactly that in a frequentist versus Bayesian environment.

We reemphasize that, while we are introducing these new precautionary parameters, we are not modifying the economic model that delivers the time-inconsistency tension, so this does not cut against Occam's Razor. The more complex dynamics we present are not obtained by expanding the economic framework; we are focused on examining parameter values for the tools employed to analyze the model, rather than parameters of the model itself. Indeed, the space of possible time series estimators is always infinite dimensional, regardless of whether a policymaker ties his hands by only considering recursive least squares. With the VALS (or "precautionary") learning we have described, the policymaker is efficiently processing the data that come out of the model given that the usual least-squares case is nested within the broader framework we consider. Doing so suggests that accounting for the variance of estimates may well assist a policymaker to escape the tragically high inflation usually expected in economies characterized by the time inconsistency of discretionary monetary policy. Indeed, precaution in learning about economic fundamentals, motivated by the potentially asymmetric nature of the policymaker's objective function and instantiated via a constant-gain VALS estimator, may cause a policymaker to target low inflation.

SUPPLEMENTARY MATERIAL

To view supplementary material for this article, please visit https://doi.org/10.1017/S1365100518000731.

NOTES

1. While we focus on a monetary policy context, models of least-squares learning are useful in a fiscal policy context as well. For example, Sargent and Wallace (1973), employing perfect foresight, find that interest rates decrease initially and then return to the steady state in response to an announced, credible, one-time increase in government spending given a balanced budget constraint. Evans et al. (2009) instead assume that agents adaptively learn about the path of future interest rates and find that announced future permanent increases in government spending can lead to low interest rates before policy implementation and then higher than steady-state interest rates after implementation before convergence to the steady state. Mitra et al. (2017) examine the size of government spending multipliers under learning versus rational expectations. Evans and Honkapohja (2003) and Bullard (2006) provide comprehensive surveys on other applications of the interaction between learning and policy.

2. See Evans and Honkapohja (1999, 2001), who provide a seminal treatment of this alternative to rational expectations.

3. The central bank's loss function can also be made explicitly asymmetric. Cukierman (2002) argues that central banks target positive inflation because their loss function is asymmetric, putting more weight on evading deflation than on evading inflation. For the dynamic consequences of an explicitly asymmetric loss function in a monetary policy context, see Ruge-Murcia (2003).

4. Cone and Shea (2017) consider differential risk postures of firms versus households when firms manage risk via the financial sector having employed learning and hedging. They find monetary policy improves welfare when learning is taxed or hedging subsidized, and that if firms are risk averse over nominal profits, then interest policies can stabilize prices thereby improving welfare.

5. We will refer to these temporary convergences to a low inflation rate as "escapes,"but the mathematics underlying them may be different from the canonical escapes of Cho et al. (2002). From Kolyuzhnov et al. (2014), we know that the approximation used by Cho et al. (2002) to characterize escape dynamics analytically breaks down at the large gains that we use for our numerical examples.

6. There is no principal-agent tension between the policy-setting division of the central bank and the econometrics division since they are both minimizing the same loss function. Willful ignorance is not entirely without theoretical precedence [see Carrillo and Mariotti (2000)]. Relatedly, within a managerial compensation context, Sinclair-Desgagne and Spaeter (2011) show that optimal incentive pay is *concave* in performance if a principal exhibits precaution over net final wealth. Here we explore the role of precaution in learning that acts as a constraint on policy making.

7. If the learning algorithm was of the decreasing gain variety, the gain would decrease with t and the economy would converge to the Nash equilibrium.

8. For robustness, we also plotted the loss surface (L) for a low value of the gain in Figure 1(b) and at a higher value in Figure 1(c) (corresponding to 400 and 10, periods respectively), provided in the Supplementary Material.

9. Those panels of Figure 4 for which the horizontal axis runs from 1 to 10, the plotted time series are the actual values for the 10,000 long simulation. Those panels for which the horizontal axis runs from 1 to 500, the plotted time series are 20-period averages. We plot 20-period average values for realized inflation and unemployment in order to smooth out the noise from the two sets of shocks.

10. The misspecification of the model will not bias estimates of the slope but will produce consistent estimates of $u_{\mathrm{NR}} + \theta x^*$ rather than u_{NR}.

11. See also Slobodyan et al. (2016) who show in a Phelps model that while a self-confirming equilibrium is stable, real-time learning diverges except for only very small gain values.

12. Since the contribution of unemployment to the loss is the same, the loss function associated with a series will move with the variability of inflation around its target value of zero. Hence we plot targeted inflation as we vary the value of the precautionary parameters over a grid.

13. Corresponding values of the slope and intercept (as in Figure 4) are available from the authors.

14. This will be the minimum average value of the loss function over long periods in a model where the public has rational expectations.

15. In the Supplementary Material, Figure 3(b) plots results given overestimates of the intercept with $a_{u_{\mathrm{NR}}} > 0$ and Figure 4(b) plots results given underestimates of the intercept with $a_{u_{\mathrm{NR}}} < 0$.

16. We thank a referee for comments leading us to the results reported in this section.

17. Carceles-Poveda and Giannitsarou (2008) and Benhabib and Dave (2014) serve as examples of this literature.

18. The optimal bias was computed assuming the model was properly specified. We have no analytic expression for the optimal bias when the ALM differs from the PLM, which is what necessitates this algorithm to learn the optimal bias.

REFERENCES

Barro, R. J. and D. B. Gordon (1983) Rules, discretion and reputation in a model of monetary policy. *Journal of Monetary Economics* 12, 101–121.

Benhabib, J. and C. Dave (2014) Learning, large deviations and rare events. *Review of Economic Dynamics* 17, 367–382.

Berger, J. O. (1985) *Statistical Decision Theory and Bayesian Analysis*, 2nd ed. New York: Springer.

Brainard, W. C. (1967) Uncertainty and the effectiveness of policy. *American Economic Review* 57, 411–425.

Bullard, J. (2006) The learnability criterion and monetary policy. *Federal Reserve Bank of St. Louis Review* 88, 203–217.

Carceles-Poveda, E. and C. Giannitsarou (2008) Asset pricing with adaptive learning. *Review of Economic Dynamics* 11, 629–651.

Carrillo, J. D. and T. Mariotti (2000) Strategic ignorance as a self-disciplining device. *The Review of Economic Studies* 67, 529–544.

Cho, I.-K. and K. Kasa (2015) Learning and model validation. *Review of Economic Studies* 82, 45–82.

Cho, I.-K., T. J. Sargent and N. Williams (2002) Escaping Nash inflation. *The Review of Economic Studies* 69, 1–40.

Cone, T. E. and P. Shea (2017) Learning, hedging, and the natural rate hypothesis. *Macroeconomic Dynamics*, 1–26. doi: http://dx.doi.org/10.1017/S1365100517000578.

Cukierman, A. (2002) Are contemporary central banks transparent about economic models and objectives and what difference does it make? *Federal Reserve Bank of St. Louis Review* 84, 15–35.

Evans, G. W. and S. Honkapohja (1999) Learning dynamics. In: J. B. Taylor and M. Woodford (eds.), *Handbook of Macroeconomics*, pp. 449-452. Amsterdam: Elsevier.

Evans, G. W. and S. Honkapohja (2001) *Learning and Expectations in Macroeconomics*. Princeton: Princeton University Press.

Evans, G. W. and S. Honkapohja (2003) Adaptive learning and monetary policy design. *Journal of Money, Credit and Banking* 35, 1045–1072.

Evans, G. W., S. Honkapohja and K. Mitra (2009) Anticipated fiscal policy and adaptive learning. *Journal of Monetary Economics* 56, 930–953.

Evans, G. W., S. Honkapohja and N. Williams (2010) Generalized stochastic gradient learning. *International Economic Review* 51, 237–262.

Kolyuzhnov, D., A. Bogomolova and S. Slobodyan (2014) Escape dynamics: A continuous-time approximation. *Journal of Economic Dynamics and Control* 38, 161–183.

Kydland, F. and E. C. Prescott (1977) Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy* 85, 473–491.

Leland, H. E. (1968) Saving and uncertainty: The precautionary demand for saving. *The Quarterly Journal of Economics* 82, 465–473.

McGough, B. (2006) Shocking escapes. *Economic Journal* 116, 507–528.

Mitra, K., G. W. Evans and S. Honkapohja (2017) Fiscal policy multipliers in an RBC model with learning. *Macroeconomic Dynamics*, 1–44. doi: http://dx.doi.org/10.1017/S1365100516001176.

Ruge-Murcia, F. J. (2003) Inflation targeting under asymmetric preferences. *Journal of Money, Credit and Banking* 35, 763–785.

Sandmo, A. (1970) The effect of uncertainty on saving decisions. *Review of Economic Studies* 37, 353–360.

Sargent, T. J. (1999) *The Conquest of American Inflation*. Princeton: Princeton University Press.

Sargent, T. J. and N. Wallace (1973) The stability of models of money and growth with perfect foresight. *Econometrica* 41, 1043–1048.

Sargent, T. J. and N. Williams (2005) Impacts of priors on convergence and escapes from Nash inflation. *Review of Economic Dynamics* 8, 360–391.

Sinclair-Desgagne, B. and S. Spaeter (2011) The Prudent Principal. SSRN Working Paper: No. 1953548.

Sims, C. A. (1988) Projecting policy effects with statistical models. *Revista de Analisis Economico* 3, 3–20.

Slobodyan, S., A. Bogomolova and D. Kolyuzhnov (2016) Stochastic gradient learning and instability: An example. *Macroeconomic Dynamics* 20, 777–790.

Tetlow, R. and P. von zur Muehlen (2004) Avoiding Nash inflation: Bayesian and robust responses to model uncertainty. *Review of Economic Dynamics* 7, 869–899.

Varian, H. R. (1975) A Bayesian approach to real estate assessment. In: S. E. Fienberg and A. Zellner (eds.), *Studies in Bayesian Econometrics and Statistics in Honor of Leonard J. Savage*. Amsterdam: North-Holland.

Zellner, A. (1986) Bayesian estimation and prediction using asymmetric loss functions. *Journal of the American Statistical Association* 81, 446–451.