

Robustness and U.S. Monetary Policy Experimentation*

Timothy Cogley[†]
Riccardo Colacito[‡]
Lars Peter Hansen[§]
Thomas J. Sargent[¶]

January 11, 2008

Abstract

We study how a concern for robustness modifies a policy maker's incentive to experiment. A policy maker has a prior over two submodels of inflation-unemployment dynamics. One submodel implies an exploitable trade-off, the other does not. Bayes' law gives the policy maker an incentive to experiment. The policy maker fears that both submodels and his prior probability distribution over them are misspecified. We compute decision rules that are robust to misspecifications of each submodel and of the prior distribution over submodels. We compare robust rules to ones that Cogley, Colacito, and Sargent (2007) computed assuming that the models and the prior distribution are correctly specified. We explain how the policy maker's desires to protect against misspecifications of the submodels, on the one hand, and misspecifications of the prior over them, on the other, have different effects on the decision rule.

KEY WORDS: Learning, model uncertainty, Bayes' law, Phillips curve, experimentation, robustness, pessimism, entropy.

*We thank Klaus Neusser and a referee for very thoughtful comments on an earlier draft.

[†]University of California, Davis. Email: twcogley@ucdavis.edu

[‡]The University of North Carolina at Chapel Hill, Kenan-Flagler Business School. Email: riccardo.colacito@unc.edu

[§]University of Chicago. Email: l-hansen@uchicago.edu

[¶]New York University and Hoover Institution. Email: ts43@nyu.edu

1 Introduction

Central bankers frequently emphasize the importance of taking parameter and/or model uncertainty into account when making decisions (e.g., see Greenspan 2004 and M. King 2004). A natural way to do so is to cast an optimal policy problem as a Bayesian decision problem. One consequence is that the decision maker’s posterior about parameters and model probabilities becomes part of the state vector. A Bellman equation would then instruct the decision maker to experiment with an eye toward tightening that posterior in the future. Although experimentation would cause near-term outcomes to deteriorate, it would speed learning and improve outcomes in the long run. Cogley, Colacito, and Sargent (2007), Wieland (2000a,b), and Beck and Wieland (2002) study aspects of this tradeoff in a variety of monetary policy models. Whether the decision maker should experiment a little or a lot depends on details of the model, but all such studies agree that an optimal policy should include some experimentation.

Despite this, prominent macroeconomists like Blinder (1998) and Lucas (1981) have forcefully recommended against purposefully experimenting on real economies in order to refine the policy maker’s knowledge (see Cogley, Colacito, and Sargent (2007) for quotations of Blinder and Lucas). An aversion to experimentation also runs through Friedman’s advocacy of a k -percent money growth rule. Resolving this ‘Bellman versus Lucas’ difference of opinion seems to require challenging the Bellman equation that leads to the recommendation to experiment purposefully.

That is what we do in this paper. In particular, we challenge the ingredient of the Bellman equation that specifies that the policy maker completely trusts his stochastic specification.¹ Our decision maker distrusts his stochastic specification and this modifies his Bellman equation. We formulate distrust by using risk-sensitivity operators and we study how that alters incentives to experiment.

As a laboratory, we adopt the model of Cogley, Colacito, and Sargent (2007). That paper computed the benefits to a Bayesian decision maker of intentional experimentation designed to reduce uncertainty about the correct model specification. The authors gave a policy maker two submodels that have very different operating characteristics that are important for policy. They also assumed that the monetary authority’s doubts are limited to not knowing the ‘correct’ value of one hyperparameter, α , the probability that one of two competing submodels generates the data. In other words, they assumed that the monetary authority has narrowed the set of possible models down to two and that it knows each submodel perfectly. If in practice one thinks that the monetary authority’s doubts are broader and vaguer,

¹Thus, Marimon (1997) noted that a Bayesian ‘knows the truth’ from the outset, so that Bayesian learning problems are just about conditioning, not constructing new joint distributions over unknowns and data.

then their calculations substantially understated the difficulty of the decision problem confronting the policy maker. For instance, the decision maker might be unsure about parameters of each submodel, might suspect that additional submodels are relevant, and might also have qualms about whether his prior adequately represents his beliefs.² The robustness calculations in this paper are designed to address some of these concerns. As we shall see, a robust decision maker still has an incentive to experiment, but the degree of experimentation is tempered by concerns that the decision problem is misspecified.

We use two risk-sensitivity operators defined by Hansen and Sargent (2005, 2007a) to construct a Bellman equation that acknowledges that the policy maker distrusts his model specification wants a decision rule that will be good enough despite the misspecification of his model. ‘Good enough’ means that a decision rule attains an acceptable outcome for a *set* of stochastic specifications centered on the policy maker’s baseline model. As we shall see, our risk-sensitivity operators summarize how the policy maker does a worst-case analysis in order to design a robust rule.

Our robust policy maker achieves a robust decision rule by pretending to be a pessimist. But pessimistic about *what*? Any worst-case analysis is context specific in the sense that ‘worst’ is relative to a particular objective function. Our decision maker attains robustness by finding a worst-case rule for a particular Kydland-Prescott (1977) ad hoc criterion for assessing inflation and unemployment outcome paths. As we vary the weights on inflation and unemployment in that criterion, what is worst changes. That affects the robust decision rule in ways that we are about to study.

1.1 Organization

Section 2 formulates a Bellman equation without concerns about misspecification. Section 3 reformulates the Bellman equation to reflect how the decision maker responds to fears that his prior over the two submodels as well as the submodels themselves are misspecified. Section 4 describes our quantitative findings. Section 5 adds some concluding remarks. We consign many technical details to an appendix.

²O’Hagan (1998, p. 22) states that “to elicit a genuine prior distribution (and typically what is needed is a joint distribution in several dimensions) is a complex business demanding a substantial effort on the part of both the statistician and the person whose prior beliefs are to be elicited.” Applied Bayesians frequently take shortcuts, such as assuming that parameters are independent a priori or choosing functional forms for convenience and not from conviction. Consequently, one might question whether a prior probability model accurately reflects the decision maker’s initial beliefs.

2 The experimentation problem without model ambiguity

A decision maker wants to maximize the following function of states s_t and controls v_t :

$$E_0 \sum_{t=0}^{\infty} \beta^t r(s_t, v_t). \quad (1)$$

The observable and unobservable components of the state vector, s_t and z_t , respectively, evolve according to a law of motion

$$s_{t+1} = \pi_s(s_t, v_t, z_t, \epsilon_{t+1}), \quad (2)$$

$$z_{t+1} = z_t, \quad (3)$$

where ϵ_{t+1} is an i.i.d. vector of shocks and $z_t \in \{1, 2\}$ is a hidden state variable that indexes submodels. Since the state variable z_t is time invariant, specification (2)-(3) states that one of the two submodels governs the data for all periods. But z_t is unknown to the decision maker. The decision maker has a prior probability $\text{Prob}(z = 1) = \alpha_0$. Where $s^t = s_t, s_{t-1}, \dots, s_0$, the decision maker recursively computes $\alpha_t = \text{Prob}(z = 1 | s^t)$ by applying Bayes' law:

$$\alpha_{t+1} = \pi_\alpha(\alpha_t, \pi_s(s_t, v_t, z_t, \epsilon_{t+1})). \quad (4)$$

Because he does not know z_t , the policy maker's prior probability α_t becomes a state variable in a Bellman equation that captures his incentive to experiment. Let asterisks denote next-period values and express the Bellman equation as

$$V(s, \alpha) = \max_v \left\{ r(s, v) + E_z [E_{s^*, \alpha^*}(\beta V(s^*, \alpha^*) | s, v, \alpha, z) | s, v, \alpha] \right\}, \quad (5)$$

subject to

$$s^* = \pi_s(s, v, z, \epsilon^*), \quad (6)$$

$$\alpha^* = \pi_\alpha(\alpha, \pi_s(s, v, z, \epsilon^*)). \quad (7)$$

E_z denotes integration with respect to the distribution of the hidden state z that indexes submodels, and E_{s^*, α^*} denotes integration with respect to the joint distribution of (s^*, α^*) conditional on (s, v, α, z) .

3 Experimentation with model ambiguity

Bellman equation (5) invites us to consider two types of misspecification of the stochastic structure: misspecification of the distribution of (s^*, α^*) conditional on

(s, v, α, z) , and misspecification of the probability α over submodels z . Following Hansen and Sargent (2005, 2007a), we introduce two risk-sensitivity operators that can help the decision maker construct a decision rule that is robust to these types of misspecification. While we refer to them as “risk-sensitivity” operators, it is actually their dual interpretations that interest us. Under these dual interpretations, a risk-sensitivity adjustment is an outcome of a minimization problem that assigns worst-case probabilities subject to a penalty on relative entropy. Thus, we view the operators as adjusting probabilities in cautious ways that assist the decision maker design robust policies.³

3.1 Two risk-sensitivity operators

3.1.1 T^1 operator

The risk-sensitivity operator T^1 helps the decision maker guard against misspecification of a submodel. Let $W(s^*, \alpha^*)$ be a measurable function of (s^*, α^*) . In our application, W will be a continuation value function. Instead of taking conditional expectations of W , we shall apply the operator:

$$T^1(W(s^*, \alpha^*))(s, \alpha, v, z; \theta_1) = -\theta_1 \log E_{s^*, \alpha^*} \exp\left(\frac{-W(s^*, \alpha^*)}{\theta_1}\right) \Big| (s, \alpha, v, z). \quad (8)$$

This operator yields the indirect utility function for a problem in which the decision maker chooses a worst-case distortion to the conditional distribution for (s^*, α^*) in order to minimize the expected value of a value function W plus an entropy penalty. That penalty limits the set of alternative models against which the decision maker guards. The size of that set is constrained by the parameter θ_1 and is decreasing in θ_1 , with $\theta_1 = +\infty$ signifying the absence of a concern for robustness. The solution to this minimization problem implies a multiplicative distortion to the Bayesian conditional distribution over (s^*, α^*) . The worst-case distortion is proportional to

$$\exp\left(\frac{-W(s^*, \alpha^*)}{\theta_1}\right), \quad (9)$$

where the factor of proportionality is chosen to make this non-negative random variable have conditional expectation equal to unity. Notice that the scaling factor and the outcome of applying the T^1 operator will depend on the state z indexing submodels even though W does not. In appendix A, we discuss in more detail a formula for this worst-case conditional distribution. Notice how (9) pessimistically twists the conditional density of (s^*, α^*) by upweighting outcomes that have lower value.

³Direct motivations for risk sensitivity can be found in Kreps and Porteus (1979) and Klibanoff, Marinacci, and Mukerji (2005).

3.1.2 T^2 operator

The risk-sensitivity operator T^2 helps the decision maker evaluate a continuation value function U that is a measurable function of (s, α, v, z) in a way that guards against misspecification of his prior α :

$$T^2(\widetilde{W}(s, \alpha, v, z))(s, \alpha, v; \theta_2) = -\theta_2 \log E_z \exp\left(\frac{-\widetilde{W}(s, \alpha, v, z)}{\theta_2}\right) \Big| (s, \alpha, v) \quad (10)$$

This operator yields the indirect utility function for a problem in which the decision maker chooses a distortion to his Bayesian prior α in order to minimize the expected value of a function $\widetilde{W}(s, \alpha, v, z)$ plus an entropy penalty. Once again, that penalty constrains the set of alternative specifications against which the decision maker wants to guard, with the size of the set decreasing in the parameter θ_2 . The worst-case distortion to the prior over z is proportional to

$$\exp\left(\frac{-\widetilde{W}(s, \alpha, v, z)}{\theta_2}\right), \quad (11)$$

where the factor of proportionality is chosen to make this nonnegative random variable have mean one. The worst-case density distorts the Bayesian probability by putting higher probability on outcomes with lower continuation values. See appendix A for more details about the worst-case density for z .⁴

Our decision maker directly distorts the date t posterior distribution over the hidden state, which in our example indexes the unknown model, subject to a penalty on relative entropy. The source of this distortion could be a change in a prior distribution at some initial date or it could be a past distortion in the state dynamics conditioned on the hidden state or model.⁵ Rather than being specific about this source of misspecification and updating all of the potential probability distributions in accordance with Bayes rule with the altered priors or likelihoods, our decision maker directly explores the impact of changes in the posterior distribution on his objective.

⁴The worst-case model as we have depicted it will depend on the endogenous state variable s_t . Since this worst-case model distorts the distribution of ϵ_{t+1} , we may prefer to represent this distortion without explicit dependence on an endogenous state variable. This can often be done for decision problems without hidden states using a ‘Big K, little k’ argument of a type featured in chapters 7 and 12 of Hansen and Sargent (2007b). A more limited notion of a worst-case model can be constructed when hidden states are present, as discussed in Hansen and Sargent (2007a).

⁵A change in the state dynamics would imply a misspecification in the evolution of the state probabilities.

3.2 A Bellman equation for inducing robust decision rules

Following Hansen and Sargent (2005, 2007a), we induce robust decision rules by replacing the mathematical expectations in (5) with risk-sensitivity operators. In particular, we substitute $(\mathbb{T}^1)(\theta_1)$ for E_{s^*, α^*} and replace E_z with $(\mathbb{T}^2)(\theta_2)$. This delivers a Bellman equation

$$V(s, \alpha) = \max_v \left\{ r(s, v) + \mathbb{T}^2 \left[\mathbb{T}^1(\beta V(s^*, \alpha^*)(s, v, \alpha, z; \theta_1)) \right] (s, v, \alpha; \theta_2) \right\}. \quad (12)$$

We find it convenient to separate the two risk-sensitivity operators by allowing for the parameters θ_1 and θ_2 to differ. The \mathbb{T}^1 operator explores the impact of *forward-looking* distortions in the state dynamics and the \mathbb{T}^2 operator explores *backward-looking* distortions in the outcome of predicting the current hidden state given current and past information. As we will see, applications of these two operators have very different ramifications for experimentation, and for that reason we find it natural to explore them separately.⁶

3.3 The submodels

Each submodel of Colacito, Cogley, and Sargent (2007) has the form

$$s_{t+1} = A_z s_t + B_z v_t + C_z \epsilon_{t+1}, \quad (13)$$

$z = 1, 2$, where s_t is an observable state vector, v_t is a control vector, and ϵ_{t+1} is an i.i.d. Gaussian processes with mean zero and contemporaneous covariance matrix I . Let $F(\cdot)$ denote the c.d.f. of this normalized multivariate Gaussian distribution. At time t , the policy maker has observed a history of outcomes $s^t = s_t, s_{t-1}, \dots, s_0$ and assigns probability α_t to model 1 and probability $(1 - \alpha_t)$ to model 2.

To capture an old debate between advocates of the natural unemployment hypothesis and those who thought that there was an exploitable unemployment-inflation trade-off, we imagine that a monetary policy authority has the following two models of inflation-unemployment dynamics:⁷

- Model 1 (Samuelson-Solow):

$$\begin{aligned} U_t &= .0023 + .7971U_{t-1} - .2761\pi_t + .0054\eta_{1,t} \\ \pi_t &= v_{t-1} + .0055\eta_{3t} \end{aligned}$$

⁶When $\theta_1 = \theta_2$ the two operators applied in conjunction give the recursive formulation of risk sensitivity proposed in Hansen and Sargent (1995), appropriately modified for the inclusion of hidden states.

⁷We use these specifications in order to have good fitting models, to keep the dimension of the state to a minimum, and still to allow ourselves to represent ‘natural rate’ and ‘non-natural rate’ theories of unemployment. For details, see appendix D of Cogley, Colacito, and Sargent.

- Model 2 (Lucas):

$$U_t = .0007 + .8468U_{t-1} - .2489(\pi_t - v_{t-1}) + .0055\eta_{2,t}$$

$$\pi_t = v_{t-1} + .0055\eta_{4t}$$

U_t is the deviation of the unemployment rate from an exogenous measure of a natural rate U_t^* , π_t is the quarterly rate of inflation, v_{t-1} is the rate of inflation that at time $t-1$ the monetary authority and private agents had both expected to prevail at time t , and, for $i = 1, 2, 3, 4$, η_{it} are i.i.d. Gaussian sequences with mean zero and variance 1. The monetary authority has a Kydland-Prescott (1977) loss function $E_0 \sum_{t=0}^{\infty} \beta^t r_t$, where $r_t = -.5(U_t^2 + \lambda v_t^2)$ and E_0 is the mathematical expectation conditioned on s_0, α_0 . The monetary authority sets v_t as a function of time t information. The analysis of Cogley, Colacito, and Sargent (2007) assumed that the monetary authority knows the parameters of each model for sure and attaches probability α_0 to model 1 and probability $1 - \alpha_0$ to model 2.⁸ Although they fit the U.S. data from 1948:3-1963:I almost equally well, these two models call for very different policies toward inflation under our loss function. Model 1, whose main features many have attributed to Samuelson and Solow (1960), has an exploitable tradeoff between v_t and subsequent levels of unemployment. Having operating characteristics advocated by Lucas (1972, 1973) and Sargent (1973), model 2 has no exploitable Phillips curve: variations in the predictable part of inflation v_t affect inflation but not unemployment. If $\alpha_0 = 0$, our decision maker should implement the trivial policy $v_t = 0$ for all t .⁹ However, if $\alpha_0 > 0$, the policy maker is willing to set $v_t \neq 0$ partly to exploit a probable inflation-unemployment tradeoff and partly to refine α .

Cogley, Colacito, and Sargent (2007) study how decision rules for this problem vary with different values of the decision maker's preference parameter λ . By comparing the decision rules from (14) with those from an associated 'anticipated utility' model, they provide a way to quantify the returns from experimentation.

4 Quantitative findings

4.1 Decision rules without robustness

As a benchmark, we first display the value function and decision rules for a version of the model without robustness (i.e., for $\theta_1 = \theta_2 = +\infty$). Figures 1 and 2 depict results for $\lambda = 0.1$ and $\beta = 0.995$, the parameters favored by Cogley, Colacito, and Sargent (2007). Notice that the value function slopes upward along

⁸As we shall see below, the T^1 operator that we use in section 4.5 allows us to analyze robustness to model perturbations that can be interpreted as coefficient uncertainty.

⁹Cogley, Colacito, and Sargent adopt a timing protocol that eliminates the inflationary bias.




Figure 1: Value function $V(U, \alpha)$ without robustness for $\lambda = 0.1$.




Figure 2: Decision rules without robustness. Black lines represent optimal experiments, and the linear red lines indicate the anticipated-utility approximations defined in Cogley, Colacito, and Sargent (2007).

the α -axis. Since α is the probability that the Samuelson-Solow model is true, the upward-sloping value function means that the policy maker is better off inhabiting a Keynesian than a classical world. That is because the Samuelson-Solow model provides a lever for controlling unemployment that the Lucas model does not. The inability to control unemployment is costly when $\lambda = 0.1$ because in that case the policy maker cares a lot about unemployment.

Also notice that for most α the decision rule for programmed inflation slopes upward along the U -axis, reflecting the countercyclical nature of policy. In addition, the policy rules are approximately linear, which signifies that there is only a modest incentive to experiment. If the connection between current actions and future beliefs were disregarded, as they are in the anticipated-utility models of Cogley, Colacito, and Sargent (2007), there would be no incentive to experiment, and the problem would reduce to a linear-quadratic dynamic program, implying linear decision rules. (The linear decision rules displayed in figure 2 are the anticipated utility decision rules.) The presence of α in the state vector breaks certainty equivalence and makes decision rules nonlinear, but in our example there is only a slight departure from linearity. Optimal monetary-policy experiments involve small, opportunistic perturbations to programmed inflation relative to anticipated-utility policies, not great leaps.

4.2 Activating T^2 only: robustness with respect to the prior

Next we activate a concern for robustness by reducing θ_2 to 0.1. We chose this value partly because it has a noticeable influence on decision rules. The left panel of figure 3 plots the worst-case distortion to α_t derived formally in appendix A, and the right panel plots a pair of decision rules for inflation v_t as a function of (U_t, α_t) . Robust decision rules are shown in red and Bayesian decision rules in gray.

A robust policy maker updates α with Bayes's theorem, then twists by increasing the probability weight on the worst-case submodel. The left panel compares the worst-case probability $\tilde{\alpha}$ with the Bayesian probability α . On the boundaries where α is 0 or 1, $\tilde{\alpha} = \alpha$. Concerns that the prior is misspecified are irrelevant when there is no model uncertainty. When α lies between 0 and 1, the worst-case model weight $\tilde{\alpha}$ is always smaller than the Bayesian update α . Since α is the probability attached to the Samuelson-Solow model, the policy maker twists by reducing his prior weight on that submodel and increasing the probability on the Lucas model. This reflects that the policy maker is worse off if the Lucas model is true because then he lacks an instrument for damping fluctuations in unemployment. Thus, it is understandable that a policy maker who cares a lot about unemployment will seek robustness by setting $\tilde{\alpha}$ less than α .

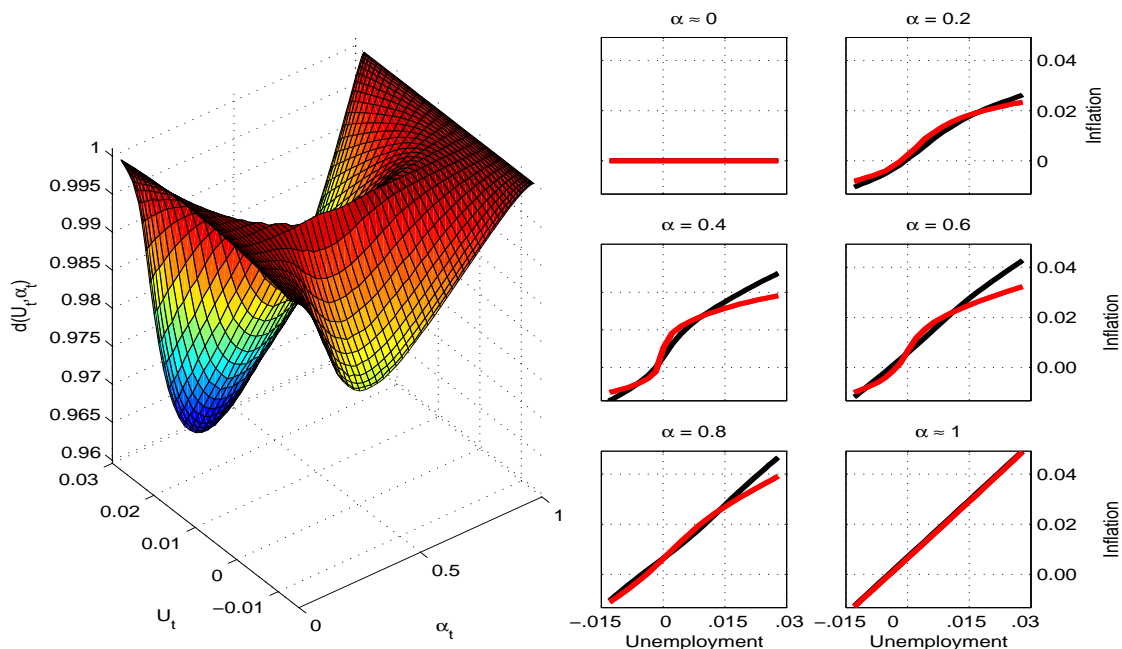


Figure 3: Robust policy with T^2 operator only, with $\lambda = 0.1$ and $\theta_2 = .1$. In the right panel, the grey line is the decision rule for the $\theta_2 = +\infty$ no-robustness decision rule, while the dark line is the $\theta_2 = 0.1$ robust decision rule.

The right-hand panel of figure 3 shows how concerns about robustness with respect to the prior over submodels alter the policy rule. Robustness matters most for intermediate values of α and high values of $|U|$. When α is close to 0 or 1, there is little model uncertainty and therefore little reason to worry about having the wrong model weight. In that case, the robust policy closely tracks the original decision rule. Similarly, the robust rule closely tracks the Bayesian policy when U is close to the point where the Samuelson-Solow model recommends zero inflation.¹⁰ In that neighborhood of U , the two models recommend similar actions, and since there is little disagreement, there is also little reason to worry about α . Robustness matters more when the models recommend very different actions, i.e. when $|U|$ is large. For intermediate values of α but high values of U , the robust decision maker sets a lower inflation target than does one who has no doubts about his prior probabilities. This is because the policy maker makes robust decisions by in effect increasing the prior weight that he attaches to the Lucas model, under which inflation is ineffective as a tool for affecting unemployment. The analysis is analogous for negative values of U ,

¹⁰This occurs when U is slightly less than zero. When $U = 0$, the Samuelson-Solow model recommends a small, positive inflation rate.

for the robust policy maker continues to twist by edging the inflation target toward zero.

Comparing these outcomes with section 5.2 of Cogley, Colacito, and Sargent (2007) shows that by expressing distrust of his prior distribution over submodels, application of the T^2 operator diminishes the incentives of the policy maker to experiment. Such distrust mutes the “opportunistic” experimentation motive that Cogley, Colacito, and Sargent (2007) found to prevail especially when $|U|$ is high.

4.3 Role of λ in determining worst-case submodel

Worst-case probabilities are context specific because they depend on the decision maker’s objective function. To bring this out, we now explore how the preceding results change as we increase the decision maker’s weight on inflation λ . A higher λ reduces the relative weight on unemployment in the period-loss function and increases the weight on inflation. Therefore, it also alters the policy maker’s perceptions about worst-case scenarios.

When λ is 16, the policy maker cares more about inflation than unemployment, and the Samuelson-Solow model becomes the worst-case scenario. Figure 4, which portrays the value function for the non-robust version of this model, shows that the value function now slopes downward along the α -axis, indicating that the authorities are better off when the Lucas model is true. When $\alpha = 0$, they refrain from varying inflation to stabilize unemployment, for they are unable to affect unemployment in any case, and they focus exclusively on maintaining price stability. That reduces inflation volatility at no cost in terms of higher unemployment volatility. Central bankers who care mostly about inflation are happier in a classical world because their job is easier in that environment.




Figure 4: Value function $V(U, \alpha)$ without robustness for $\lambda = 16$.

Figure 5 illustrates how this affects robust policies. If we were to hold θ_2 constant while increasing λ , a concern for robustness would vanish for $\lambda = 16$, so we also reduce θ_2 to 0.001 to compensate.¹¹ Because the Samuelson-Solow model is the worst case, a robust planner twists by increasing its probability weight. This explains why in the left panel the twisted model weight $\tilde{\alpha}$ is greater than α in almost all states of the world. This raises programmed inflation when unemployment is low, but because λ is so high, v_t always remains close to zero, with or without robustness. Thus, differences in the policy functions are slight, amounting to just a few basis points.

¹¹When $\theta_2 = 0.001$ and $\lambda = 0.1$, the robust planner becomes hypervigilant, and the value function ceases to be concave in v_t and convex in the choice of the perturbation to the approximating model. Whittle (1990) describes a breakdown value of θ_1 as a point of ‘utter psychotic despair’.

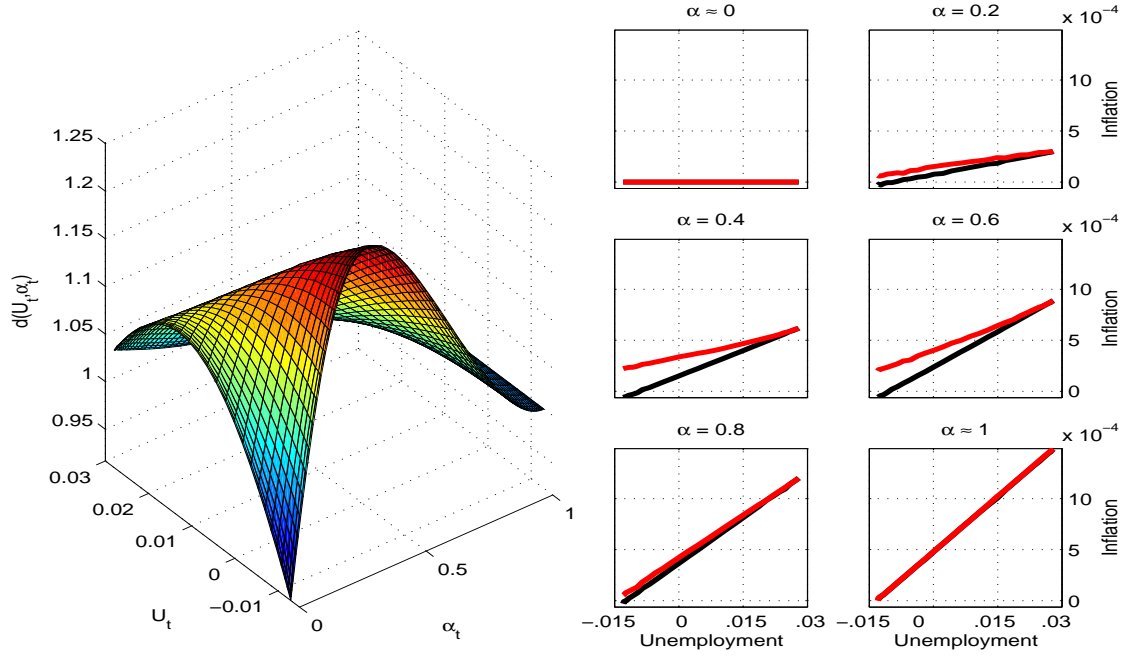


Figure 5: Robust policy with T^2 operator only, with $\lambda = 16$ and $\theta_2 = 0.001$. In the right panel, the dark line is the decision rule for the $\theta_2 = +\infty$ no-robustness decision rule, while the grey line is the $\theta_2 = 0.001$ robust decision rule.

For intermediate values of λ , either model could be the worst, so the distortion to α could go either way. It follows that a concern for robustness could make policy more or less countercyclical. For example, figures 6 and 7 display the value function, α -distortion, and decision rules, respectively, for $\lambda = 1$ and $\theta_2 = 0.001$. When inflation and unemployment are equally weighted, the non-robust value function still slopes upward, which means that the Lucas model is still associated with the worst-case scenario, and the robust planner twists in most states of the world by reducing $\tilde{\alpha}$ relative to the Bayesian update α .¹² Accordingly, the robust policy rule is still less countercyclical than the Bayesian decision rule.




Figure 6: Value function $V(U, \alpha)$ without robustness for $\lambda = 1$.

¹²An exception occurs when α is close to zero, where the robust planner twists toward the Samuelson-Solow model. This matters only slightly for policy because programmed inflation is always close to zero when α is close to zero.

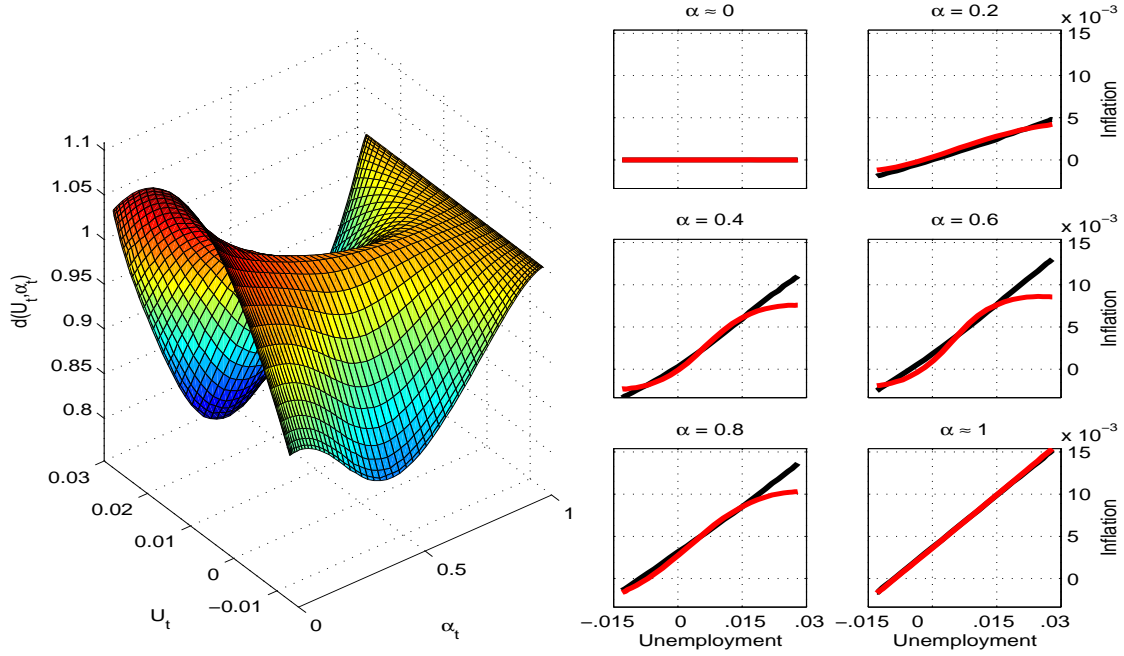


Figure 7: Robust policy with T^2 operator only, with $\lambda = 1$ and $\theta_2 = 0.001$. In the right panel, the grey line is the decision rule for the $\theta_2 = +\infty$ no-robustness decision rule, while the dark line is the $\theta_2 = 0.001$ robust decision rule.

4.4 Dwindling effect of T^2 operator

Colacito, Cogley, and Sargent (2007) indicate how, when one of the two submodels is true, α_t converges either to zero or one as $t \rightarrow +\infty$. Furthermore, even when the data are generated by a third submodel not considered by the decision maker, it is often the case that α_t still converges to zero or one. The preceding figures indicate that at the dogmatic boundary $\alpha = 0$ or $\alpha = 1$, there is no room for the T^2 operator to distort beliefs. This means that the inexorable working of Bayes' law causes the effects of the T^2 operator to die off over time.

4.5 Activating T^1 : robustness with respect to each sub-model

Keeping $\lambda = 0.1$, we now use the T^1 operator to express a concern about misspecification of the unemployment-inflation dynamics within each of the two submodels. To begin, by setting $\theta_2 = +\infty$ we shall assume that the decision maker is confident about his prior. We express a concern for misspecification of the submodels by replacing the $E_{s^*\alpha^*}$ in (14) with the T^1 operator. In particular, we replace

$$\beta \int V(A_z s_t + B_z v_t + C_z \epsilon_{t+1}, \pi_\alpha(\alpha_t, A_z s_t + B_z v_t + C_z \epsilon_{t+1})) dF(\epsilon_{t+1})$$

in (14) with $(T^1(\beta V))(s_t, \alpha_t, v_t, z; \theta_1)$.

Figures 8 and 9 display the conditional means and variances of the worst-case conditional densities (22) for the Samuelson-Solow and Lucas models, respectively, for $\theta_1 = 0.1$ and $\theta_2 = +\infty$ ¹³. The nature of the worst-case scenario is similar in the two models. In both cases, the worst-case model envisions a higher probability of drawing a deviation-amplifying shock when $|U|$ is already large. The expected values of the distorted unemployment shocks in the two models, η_1 and η_2 , are positive when unemployment is high and negative when it is low, and this directly amplifies unemployment volatility. Similarly, the expected values of the distorted inflation shocks, η_3 and η_4 , are negative when U is high and positive when U is low. This indirectly increases unemployment volatility because U varies inversely with respect to unexpected inflation. In addition, the shock variances are altered to increase volatility, being greater when $|U|$ is large.

Figure 10 displays the corresponding robust decision rule. To offset the greater risk of a deviation-amplifying shock, the robust policy authority adopts a more aggressive countercyclical stance relative to that set by a policy maker who fully trusts the specification of each model. Thus, concerns about possible misspecifications of the submodels have an opposite effect from a concern about the prior alone that we summarized in figure 3.

¹³We set $\theta_1 = .1$ because this value delivers noticeable effects on the decision rule.

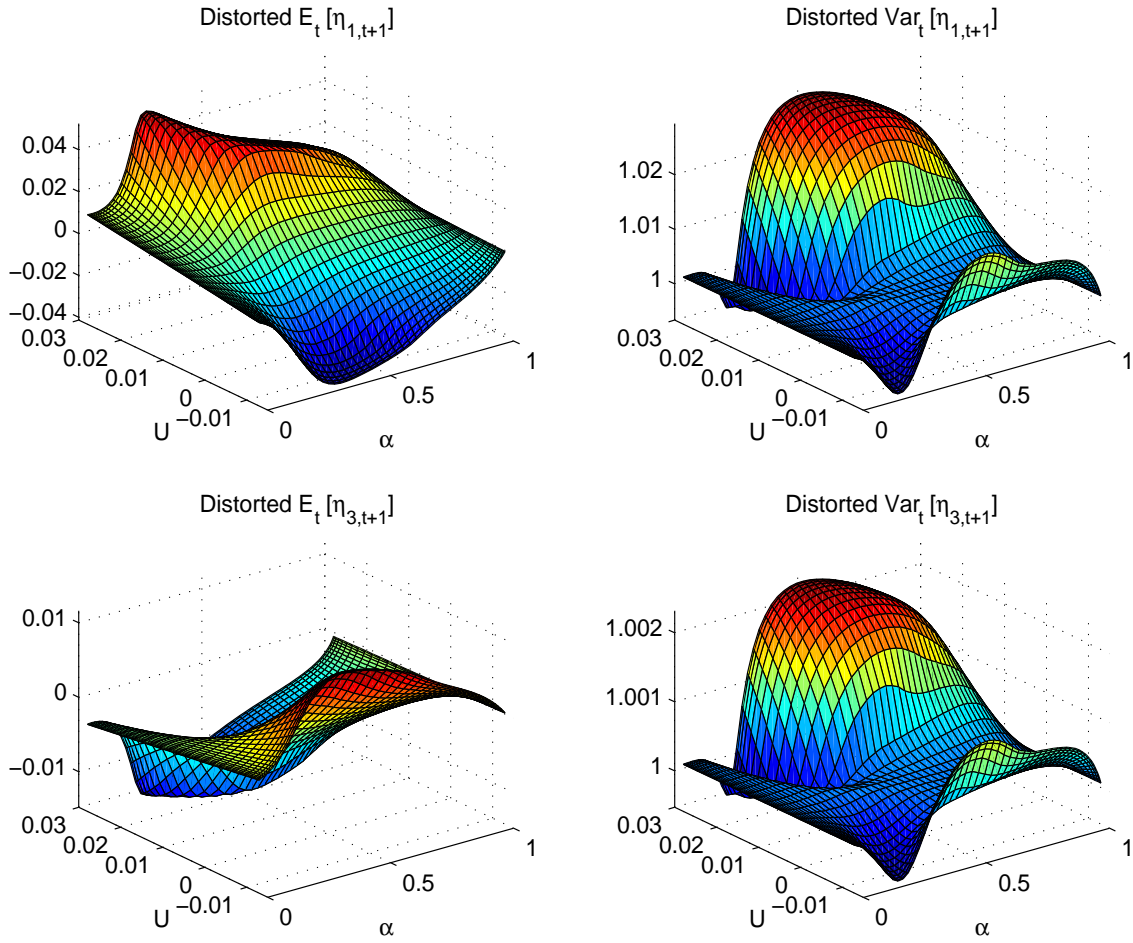


Figure 8: Conditional means and variances of distorted shocks to the Samuelson-Solow model with $\theta_1 = .1$.

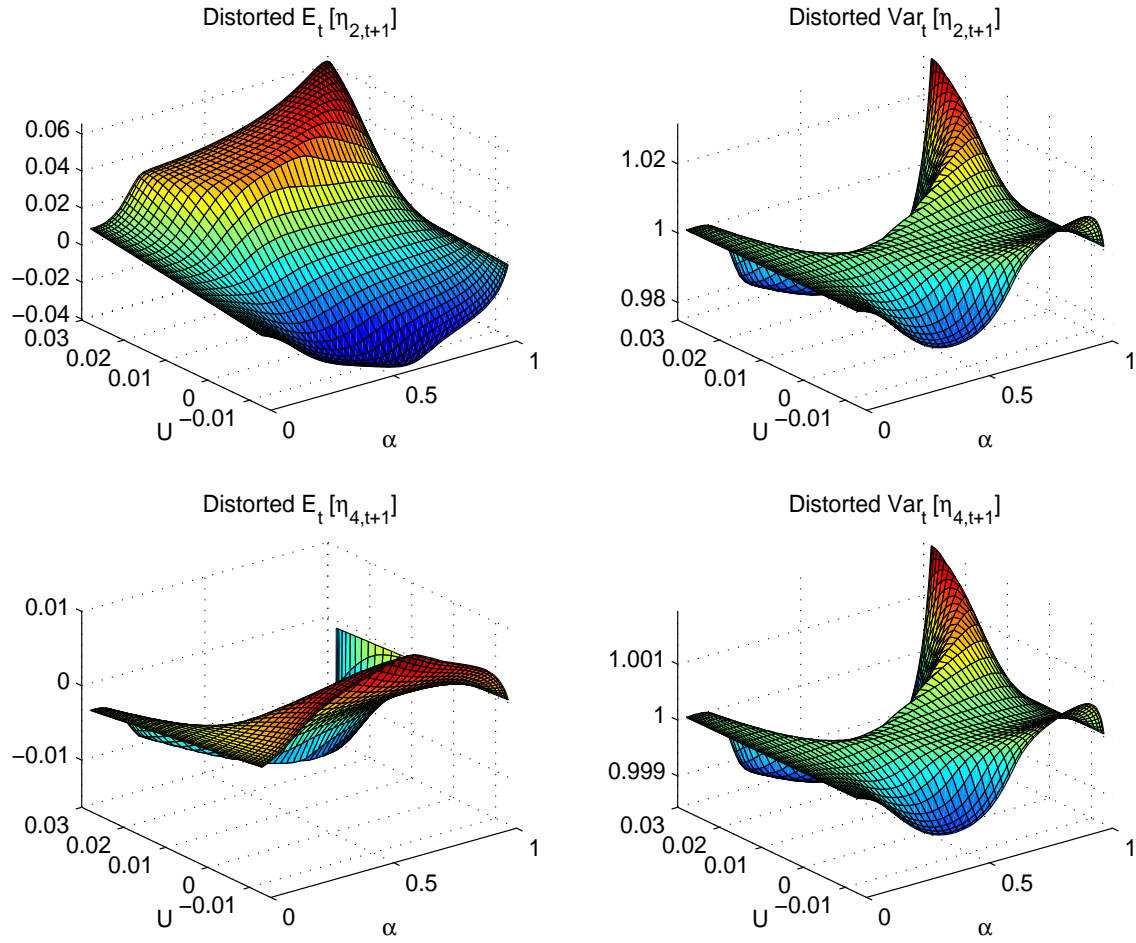


Figure 9: Conditional means and variances of distorted shocks to the Lucas model with $\theta_1 = .1$.




Figure 10: Robust policy with T^1 operator only, $\theta_1 = .1$. The grey line is the decision rule for the $\theta_1 = +\infty$ no-robustness decision rule, while the dark line is the $\theta_1 = 0.1$ robust decision rule.

4.6 How the two forms of misspecification interact: activating both the T^1 and T^2 operators

Figure 11 activates concerns about both features of the specification. As might be guessed from the complexion of the earlier results, turning on both sources of concern about robustness yields a decision rule that is close to the one we obtained without any concerns about robustness. When $\lambda = 0.1$, the decision maker makes programmed inflation less countercyclical to guard against misspecification of the prior, but makes v_t more countercyclical to protect against misspecification of the two submodels. In effect, the worst-case α shown in the left panel of figure 11 offsets the worst-case dynamics coming from the dependence of the worst-case conditional mean on (U_t, α_t) , so that the combined effects of T^1 and T^2 approximately cancel. Thus, the optimal Bayesian decision rule with experimentation – calculated without explicit reference to robustness – is robust to a mixture of concerns about the two types of misspecification.¹⁴

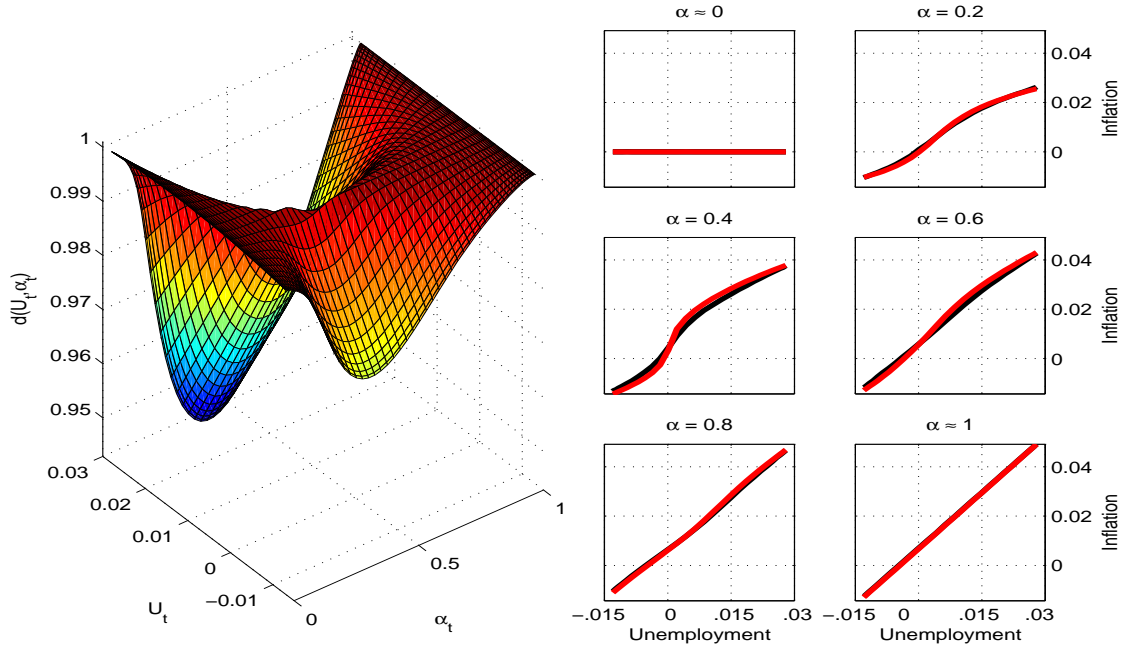


Figure 11: Worst-case α and decision rule with concerns about both source of misspecification, captured by T^1 and T^2 with $\theta_1 = \theta_2 = .1$. The black line on the right panel indicates the $\theta_1 = \theta_2 = +\infty$ decision rule and the grey line indicates the $\theta_1 = \theta_2 = .1$ decision rule.

¹⁴Results like this also obtain for other values of λ .

4.7 How the two operators influence experimentation

Next we examine more closely how the two risk-sensitivity operators affect motives to experiment. Figure 12 compares robust, Bayesian, and anticipated-utility decision rules. To highlight their differences, we set $\alpha = 0.4$ to focus on a part of the state space where experimental motives are strongest. We interpret differences of robust decision rules relative to the non-experimental, anticipated-utility decision rule. When a risk-sensitivity operator moves a decision rule closer to the anticipated-utility policy, we say that it tempers experimentation.




Figure 12: Robust, Bayesian, and Anticipated-Utility Policy Rules

The left panel of figure 12 illustrates the influence of the backward-looking T^2 operator. On balance, T^2 mutes experimentation. For small values of $|U|$, the robust and Bayesian policies are essentially the same, while for larger values the robust policy curls back toward the nonexperimental decision rule. Since the robust rule calls for no more experimentation than the Bayesian policy when unemployment is close to the natural rate and calls for less when the unemployment gap is large in magnitude, less experimentation occurs along a learning transition path.¹⁵

The middle panel examines the influence of the forward-looking T^1 operator. In this case, experimentation is muted for small values of $|U|$ but strongly enhanced for

¹⁵It is conceivable that T^2 results in more experimentation for larger values of $|U|$ not shown on the graph, but those states are rarely visited.

large values. Since realizations of $|U|$ in the neighborhood of 0.02 are not unusual,¹ typically results in more experimentation.

Finally, the right-hand panel illustrates what happens when both operators are active. Since one operator mutes and the other enhances experimentation, the two operators offset, so that a decision rule coming from the application of both operators is about the same as the Bayesian decision rule. In summary, the net effect of robustness on experimentation is ambiguous and depends on the penalty parameters θ_1 and θ_2 .

4.8 How long does it take to learn the truth?

In this subsection we analyze how a preference for robustness affects the number of quarters that are needed to learn the true model. We address this question by simulation. For each simulation, we assume that either the Lucas' model or the Samuelson and Solow model is the actual data generating process. We initialize the state space at various levels of (α_0, U_0) and let the system run according to the dynamics of the true model under the optimal inflation rate impelled by the relevant Bellman equation. For each experiment we report the median number of quarters, that are needed for α to get within a 0.01 neighborhood of what it should be under the true data generating process. We also report the 10%-90% confidence interval for each case. Each experiment is based on 1000 simulations of length 700 quarters. The results are reported in table 1, and they can be summarized by comparison to the baseline case (i.e. $(\lambda = 0.1, \theta_1 = +\infty, \theta_2 = +\infty)$).

1. A fear for prior misspecification (i.e., $\theta_1 = +\infty, \theta_2 = 0.1$) increases the time needed to learn the true model. This is particularly apparent for high initial levels of unemployment. In these cases, the distorted probability distribution makes the Samuelson and Solow model less likely. Hence, the optimal decision rule calls for a lower inflation rate, which damps experimentation and makes it harder to discover the actual data generating process (see figure 3).
2. A fear of misspecification of the probability distribution within each sub-model (i.e., $\theta_1 = 0.1, \theta_2 = +\infty$) increases the speed of convergence. In this case, the policy maker sets higher rates of inflation in the rise of unemployment (see figure 8). The higher the degree of experimentation results in a usually very quick convergence to the true model.
3. When both risk-sensitivity operators are turned on, there is no significant difference with respect to the baseline model. This is the result of the two fears of misspecification offsetting each other in the choice of the optimal inflation rate (see figure 11).

True Model	α_0	U_0	Waiting time			
			$\theta_1 = +\infty$	$\theta_1 = +\infty$	$\theta_1 = 0.1$	$\theta_1 = 0.1$
			$\theta_2 = +\infty$	$\theta_2 = 0.1$	$\theta_2 = +\infty$	$\theta_2 = 0.1$
SS	0.01	0	218	223	214	231
			[128,439]	[130,444]	[123,416]	[128,451]
SS	0.01	0.025	229	233	224	241
			[141,459]	[135,453]	[143,404]	[147,441]
Lucas	0.99	0	89	96	66	89
			[32,197]	[39,204]	[22,150]	[34,202]
Lucas	0.99	0.025	66	80	54	71
			[20,175]	[27,203]	[4,149]	[19,194]
SS	0.5	0	40	37	38	37
			[21,79]	[20,75]	[18,71]	[20,70]
SS	0.5	0.025	22	27	9	20
			[5,61]	[10,73]	[2,51]	[4,53]
Lucas	0.5	0	72	78	60	74
			[26,192]	[26, 180]	[17,160]	[23,188]
Lucas	0.5	0.025	58	71	28	59
			[15,167]	[23,187]	[2,127]	[14,177]

Table 1: Waiting Times (in quarters) for various data-generating processes and initial (α_0, U_0) pairs. The variable that we call waiting time represents the number of quarters that are needed for α to return to within a 0.01 neighborhood of what it should be under the data generating process. For each experiment, we report the true model, the initial prior, the initial unemployment rate, the median waiting time, and the 10% - 90% confidence sets in square brackets for various pairs of (θ_1, θ_2) .

5 Concluding remarks

In this paper, we study how concerns for robustness modify incentives to experiment. We use a decision theory that explores robustness of decision rules by calculating bounds on value functions over a set of probability models near a decision maker’s approximating model. Seeking bounds on value functions over a set of probability models automatically leads to a worst-case analysis. We study a setting in which a decision maker’s approximating model is an average of two submodels. The decision maker uses Bayes’ law to update priors over submodels as new data arrive. Our T^1 operator checks robustness of each submodel. Our T^2 operator checks robustness with respect to a prior over submodels.

Our working example is the model in Colacito, Cogley, and Sargent (2007) in which a Samuelson-Solow submodel offers a permanently exploitable trade-off between inflation and unemployment and another Lucas submodel lacks a trade-off that is even temporarily exploitable. This is a good setting for illustrating how the worst-case model is worst relative to the decision maker’s objective. When the monetary policy decision maker puts more weight on unemployment ($\lambda = 0.1$), the Lucas model is worse. That makes the robust policy less countercyclical than the policy that completely trusts the model. When more weight is on inflation ($\lambda = 16$), the Samuelson-Solow model is worse for the policy maker. That makes the robust policy more countercyclical than the nonrobust policy.

Robust policy makers have an incentive to experiment for the same reason that Bayesian policy makers do. The decision maker’s posterior is still an element of the state vector, so robust Bellman equations continue to instruct the decision maker to experiment with an eye toward tightening the posterior in the future. What changes are the costs and benefits of experimentation. How decision rules are altered is model specific, so robustness could in principle either enhance or mute experimentation.

In the present context, the T^1 and T^2 operators have countervailing effects on policy. When $\lambda = 0.1$, concerns that the submodels are misspecified make policy more countercyclical than in a Bayesian setting, while concerns that the prior is misspecified make policy less countercyclical. When these results are compared to Cogley, Colacito, and Sargent’s (2007) measures of the contribution of an experimentation motive to the policy rule, they show that with complete trust in the prior over submodels, distrust of the submodels increases the motive to experiment, while with complete trust in the submodels, distrust of the prior over submodels diminishes the motive to experiment. When both operators are active, their effects approximately cancel, and the robust policy well approximates the Bayesian decision rule. Since Cogley, Colacito, and Sargent (2007) showed that the Bayesian decision rule well approximates an ‘anticipated-utility’ policy that suppresses experimentation altogether, it follows that with both of our T operators active, the optimally

robust policy has little or no experimentation. Thus, the disagreement between the ordinary Bellman equation's recommendation to experiment and Blinder's and Lucas's advice not to experiment, cited at the beginning of this paper, can in principle be rationalized by using the \mathbb{T}^2 operator to express a distrust of the decision maker's prior over the submodels that offsets other motives to experiment.

A Details

A.1 The Bellman equation

Our Bellman equation without fear of model misspecification is:

$$\begin{aligned} V(s_t, \alpha_t) = & \max_{v_t} \{r(s_t, v_t) \\ & + \beta \alpha_t \int V(A_1 s_t + B_1 v_t + C_1 \epsilon_{t+1}, \pi_\alpha(\alpha_t, A_1 s_t + B_1 v_t + C_1 \epsilon_{t+1})) dF(\epsilon_{t+1}) \\ & + \beta(1 - \alpha_t) \int V(A_2 s_t + B_2 v_t + C_2 \epsilon_{t+1}, \pi_\alpha(\alpha_t, A_2 s_t + B_2 v_t + C_2 \epsilon_{t+1})) dF(\epsilon_{t+1})\} \end{aligned} \quad (14)$$

The optimal decision rule can be represented recursively as

$$v_t = v(s_t, \alpha_t). \quad (15)$$

Repeated substitution of (7) into (15) yields the policy maker's strategy in the form of a sequence of functions

$$v_t = \sigma_t(s^t, \alpha_0). \quad (16)$$

Cogley, Colacito, and Sargent (2007) derive the function $\pi_\alpha(s_t, \alpha_t)$. To summarize their calculations, let $\Omega_i = C_i C_i'$, $R_t = \frac{\alpha_t}{1 - \alpha_t}$, and define

$$\begin{aligned} g(\epsilon_{t+1}; s_t, \alpha_t) = & \log R_t - \frac{1}{2} \log |\Omega_1| + \frac{1}{2} \log |\Omega_2| - \frac{1}{2} (C_1 \epsilon_{t+1})' \Omega_1^{-1} (C_1 \epsilon_{t+1}) \\ & + \frac{1}{2} [(A_1 - A_2)s_t + (B_1 - B_2)v_t + C_1 \epsilon_{t+1}]' \\ & \times \Omega_2^{-1} [(A_1 - A_2)s_t + (B_1 - B_2)v_t + C_1 \epsilon_{t+1}] \end{aligned} \quad (17)$$

and

$$\begin{aligned} h(\epsilon_{t+1}; s_t, \alpha_t) = & \log R_t - \frac{1}{2} \log |\Omega_1| + \frac{1}{2} \log |\Omega_2| + \frac{1}{2} (C_2 \epsilon_{t+1})' \Omega_2^{-1} (C_2 \epsilon_{t+1}) \\ & - \frac{1}{2} [(A_2 - A_1)s_t + (B_2 - B_1)v_t + C_2 \epsilon_{t+1}]' \\ & \times \Omega_1^{-1} [(A_2 - A_1)s_t + (B_2 - B_1)v_t + C_2 \epsilon_{t+1}]. \end{aligned} \quad (18)$$

The Bellman equation (14) becomes

$$V(s_t, \alpha_t) = \max_{v_t} \left\{ r(s_t, v_t) + \beta \alpha_t \int V \left(A_1 s_t + B_1 v_t + C_1 \epsilon_{t+1}, \frac{e^{g(\epsilon_{t+1})}}{1 + e^{g(\epsilon_{t+1})}} \right) dF(\epsilon_{t+1}) \right. \\ \left. + \beta(1 - \alpha_t) \int V \left(A_2 s_t + B_2 v_t + C_2 \epsilon_{t+1}, \frac{e^{h(\epsilon_{t+1})}}{1 + e^{h(\epsilon_{t+1})}} \right) dF(\epsilon_{t+1}) \right\}. \quad (19)$$

Cogley, Colacito, and Sargent (2007) also describe how to approximate the solution of (19) and the robust counterpart to it that we propose in subsection 3.2.

A.2 The two operators

We describe details of how the operators $\mathbb{T}^1, \mathbb{T}^2$ apply in our particular setting.

A.3 \mathbb{T}^1 operator

For a given value function $V(s_{t+1}, \alpha_{t+1})$ and a given decision rule $v_t = v(s_t, \alpha_t)$, define

$$\begin{aligned} \mathbb{T}^1(V(s_{t+1}, \alpha_{t+1}))((s_t, \alpha_t, v_t, z; \theta_1) & \quad (20) \\ &= -\theta_1 \log \int \exp \left(-\frac{V(A_z s_t + B_z v_t + C_z \epsilon_{t+1}, \pi_\alpha(\alpha_t, A_z s_t + B_z v_t + C_z \epsilon_{t+1}))}{\theta_1} \right) dF(\epsilon_{t+1}). \\ &= \min_{\phi(s_t, v_t, \alpha_t, \epsilon_{t+1}) \geq 0} \int \left[V(A_z s_t + B_z v_t + C_z \epsilon_{t+1}, \pi_\alpha(\alpha_t, A_z s_t + B_z v_t + C_z \epsilon_{t+1})) \right. \\ & \quad \left. + \theta_1 \log \phi(s_t, v_t, \alpha_t, \epsilon_{t+1}) \right] \phi(s_t, v_t, \alpha_t, \epsilon_{t+1}) dF(\epsilon_{t+1}) \end{aligned} \quad (21)$$

where the minimization is subject to $E[\phi(s_t, v_t, \alpha_t, \epsilon_{t+1}) | s_t, \alpha_t, v_t, j] = 1$. The minimizer in (21) is a worst-case distortion to the density of ϵ_{t+1} :

$$\phi^*(\epsilon_{t+1}, s_t, \alpha_t) = \frac{\exp \left(-\frac{V(A_z s_t + B_z v_t + C_z \epsilon_{t+1}, \pi_\alpha(\alpha_t, A_z s_t + B_z v_t + C_z \epsilon_{t+1}))}{\theta_1} \right)}{\int \exp \left(-\frac{V(A_z s_t + B_z v_t + C_z \epsilon_{t+1}, \pi_\alpha(\alpha_t, A_z s_t + B_z v_t + C_z \epsilon_{t+1}))}{\theta_1} \right) dF(\epsilon_{t+1})},$$

where it is understood that v_t on the right side is evaluated at a particular decision rule $v(s_t, \alpha_t)$. The distorted conditional density of ϵ_{t+1} is then

$$\check{\phi}(\epsilon_{t+1}, s_t, \alpha_t) = \phi_n(\epsilon_{t+1}) \phi^*(\epsilon_{t+1}, s_t, \alpha_t), \quad (22)$$

where $\phi_n(\epsilon_{t+1})$ is the standard normal density.

A.4 \mathbb{T}^2 operator

For $j = 1, 2$, let $\check{V}(s, \alpha, v, z)$ be distinct functions of (s, α, v) for $z = 0, 1$. Define

$$\begin{aligned} \mathbb{T}^2(\check{V}(s, \alpha, v, z; \theta_2))(s, \alpha, v) & \quad (23) \\ &= -\theta_2 \log \left[\alpha \exp\left(\frac{-\check{V}(s, \alpha, v, 0)}{\theta_2}\right) + (1 - \alpha) \exp\left(\frac{-\check{V}(s, v, \alpha, 1)}{\theta_2}\right) \right] \\ &= \min_{\psi_0 \geq 0, \psi_1 \geq 0} \left\{ [\check{V}(s, \alpha, v, 0) + \theta_2 \log \psi_1] \psi_0 \alpha \right. \\ &\quad \left. + [\check{V}(s, \alpha, v, 1) + \theta_2 \log \psi_1] \psi_1 (1 - \alpha) \right\} \quad (24) \end{aligned}$$

where the minimization is subject to $\psi_0 \alpha + \psi_1 (1 - \alpha) = 1$. The minimizers of (24) are

$$\begin{aligned} \psi_0^*(s, \alpha, v) &= k \exp\left(\frac{-\check{V}(s, v, \alpha, 0)}{\theta_2}\right) \\ \psi_1^*(s, \alpha, v) &= k \exp\left(\frac{-\check{V}(s, v, \alpha, 1)}{\theta_2}\right) \end{aligned}$$

where $k^{-1} = \exp\left(\frac{-\check{V}(s, \alpha, v, 0)}{\theta_2}\right) \alpha + \exp\left(\frac{-\check{V}(s, v, \alpha, 1)}{\theta_2}\right) (1 - \alpha)$. An associated worst-case probability that $z = 0$ is given by

$$\hat{\alpha} = \psi_0^*(s, \alpha, v) \alpha. \quad (25)$$

References

- Beck, G. and V. Wieland, 2002, “Learning and Control in a Changing Environment,” *Journal of Economic Dynamics and Control* 26, 1359-1377.
- Blinder, A.S. 1998, *Central Banking in Theory and Practice*, Cambridge: MIT Press.
- Cogley, T., R. Colacito, and T.J. Sargent (2007), “Benefits from U.S. Monetary Policy Experimentation in the Days of Samuelson and Solow and Lucas,” *Journal of Money, Credit and Banking*, supplement to Vol. 39, pp. 67–99.
- Greenspan, A., 2004, “Risk and Uncertainty in Monetary Policy,” *American Economic Review Papers and Proceedings* 94(2), 33-40.
- Hansen, Lars Peter and Thomas J. Sargent, 1995, “Discounted Linear Exponential Quadratic Gaussian Control,” *IEEE Transactions on Automatic Control* 40(5), 968-971.

- Hansen, Lars Peter and Thomas J. Sargent, 2005, "Robust Estimation and Control under Commitment," *Journal of Economic Theory*, Vol. 124, No. 2, pp. 258–301.
- Hansen, Lars Peter and Thomas J. Sargent, 2007a, "Robust Estimation and Control without Commitment," *Journal of Economic Theory*, 136(1), 1-27.
- Hansen, Lars Peter and Thomas J. Sargent, 2007b. *Robustness*. Princeton University Press, Princeton New Jersey.
- King, M., 2004, "Innovations and Issues in Monetary Policy: Panel Discussion," *American Economic Review Papers and Proceedings* 94(2), 43-45.
- Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji, 2005, "A Smooth Model of Decision Making Under Uncertainty," *Econometrica*, 73, 1840-1892.
- Kreps, David M. and Evan L. Porteus, 1978, "Temporal Resolution of Uncertainty and Dynamic Choice," *Econometrica*, 46, 185-200.
- Kydland, F.E. and E.C. Prescott, 1977, "Rules Rather Than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy* 85, 473-91.
- Lucas, R.E., Jr., 1972, "Expectations and the Neutrality of Money," *Journal of Economic Theory* 4, 103-124.
- Lucas, R.E., Jr., 1973, "Some International Evidence on Output-Inflation Trade-Offs," *American Economic Review* 63, 326-334.
- Lucas, R.E., Jr., 1981, "Methods and Problems in Business Cycle Theory," in R.E. Lucas, Jr., (ed.) *Studies in Business-Cycle Theory*, Cambridge, Mass., MIT Press.
- Marimon, Ramon, 1997, "Learning from Learning in Economics," in David Kreps and Kenneth Wallis (eds.), *Advances in Economics and Econometrics: Theory and Applications. Seventh World Congress, Volume 1.* Cambridge University Press.
- O'Hagan, A., 1998, "Eliciting Expert Beliefs in Substantial Practical Applications," *The Statistician* 47, Part 1, 21-35.
- Samuelson, P.A., and R.M. Solow, 1960, "Analytical Aspects of Anti-Inflation Policy," *American Economic Review*, 50, 177–184.
- Sargent, T.J., 1973, "Rational Expectations, the Real Rate of Interest, and the Natural Rate of Unemployment," *Brookings Papers on Economic Activity*, 429-472.
- Whittle, Peter, 1990, *Risk-Sensitive Optimal Control*, John Wiley & Sons, New York.
- Wieland, V., 2000a, "Monetary Policy, Parameter Uncertainty, and Optimal Learning," *Journal of Monetary Economics* 46, 199-228.
- Wieland, V., 2000b, "Learning by Doing and the Value of Optimal Experimentation," *Journal of Economic Dynamics and Control* 24, 501-534.