

WANTING ROBUSTNESS IN MACROECONOMICS

LARS PETER HANSEN AND THOMAS J. SARGENT

1. INTRODUCTION

1.1. **Foundations.** von Neumann and Morgenstern (1944), Savage (1954), and Muth (1961) created mathematical foundations that applied economists have used to construct quantitative dynamic models for policy making. These foundations give modern dynamic models an internal coherence that lead to sharp empirical predictions. When we acknowledge that models are approximations, logical problems emerge that unsettle those foundations. Because the rational expectations assumption works the presumption of a correct specification particularly hard, admitting model misspecification raises especially interesting problems about how to extend rational expectations models.¹

A model is a probability distribution over a sequence. The rational expectations hypothesis delivers empirical power by imposing a communism of models: the people being modeled, the econometrician, and nature share the same model, *i.e.*, the same probability distribution over sequences of outcomes. This ‘communism’ is used both in solving a rational expectations model and when a law of large numbers is appealed to when justifying GMM or maximum likelihood estimation of model parameters. Imposition of a common model removes economic agents’ models as objects that require separate specification. The rational expectations hypothesis converts agents’ beliefs from model *inputs* to model *outputs*.

The idea that models are approximations puts more models in play than the rational expectations equilibrium concept handles. To say that a model is an approximation is to say that it approximates another model. Viewing models as approximations requires somehow reforming the common models requirement imposed by rational expectations.

The consistency of models imposed by rational expectations has profound implications about the design and impact of macroeconomic policy-making, *e.g.* see Lucas (1976) and Sargent and Wallace (1975). There is relatively little work studying how

Date: May 28, 2010.

We thank Ignacio Presno, Robert Tetlow, François Velde, Neng Wang, and Michael Woodford for insightful comments on earlier drafts.

¹Applied dynamic economists readily accept that their models are tractable approximations. Sometimes we express this by saying that our models are abstractions or idealizations. Other times we convey it by focusing a model only on ‘stylized facts’.

those implications would be modified within a setting that explicitly acknowledges decision makers' fear of model misspecification.²

Thus, the idea that models are approximations conflicts with the von Neumann-Morgenstern-Savage foundations for expected utility and with the supplementary equilibrium concept of rational expectations that underpins modern dynamic models. In view of those foundations, treating models as approximations raises three questions. What standards should be imposed when *testing* or *evaluating* dynamic models? How should private decision makers be modeled? How should macroeconomic policy-makers use misspecified models? This essay focuses primarily on the latter two questions. But in addressing these questions we are compelled to say something about testing and evaluation.

This essay describes an approach in the same spirit but differing in many details from Epstein and Wang (1994). We follow Epstein and Wang in using the Ellsberg paradox to motivate a decision theory for dynamic contexts that is based on the min-max theory with multiple priors of Gilboa and Schmeidler (1989). We differ from Epstein and Wang (1994) in drawing our formal models from recent work in control theory. This choice leads to many interesting technical differences in the particular class of models against which our decision maker prefers robust decisions. Like Epstein and Wang (1994), we are intrigued by a passage from Keynes (1936):

A conventional valuation which is established as the outcome of the mass psychology of a large number of ignorant individuals is liable to change violently as the result of a sudden fluctuation in opinion due to factors which do not really make much difference to the prospective yield; since there will be no strong roots of conviction to hold it steady.

Epstein and Wang provide a model of asset price indeterminacy that might explain the sudden fluctuations in opinion that Keynes mentions. In Hansen and Sargent (2008a), we offer a model of sudden fluctuations in opinion coming from a representative agent's difficulty in distinguishing between two models of consumption growth that differ mainly in their implications about hard-to-detect low frequency components of consumption growth. We describe this force for sudden changes in beliefs in section 5.5 below.

2. KNIGHT, SAVAGE, ELLSBERG, GILBOA-SCHMEIDLER, AND FRIEDMAN

In *Risk, Uncertainty and Profit*, Frank Knight (1921) envisioned profit-hunting entrepreneurs who confront a form of uncertainty not captured by a probability model.³

²But see Karantounias (2009), Woodford (2010), Hansen and Sargent (2008b), chapters 15 and 16, and Orlik and Presno (2009).

³See Epstein and Wang (1994) for a discussion containing many of the ideas summarized here.

He distinguished between risk and uncertainty, and reserved the term risk for ventures with outcomes described by known probabilities. Knight thought that probabilities of returns are not known for many physical investment decisions. Knight used the term uncertainty to refer to such unknown outcomes.

After Knight (1921), Savage (1954) contributed an axiomatic treatment of decision-making in which preferences over gambles could be represented by maximizing expected utility under subjective probabilities. Savage's work extended the earlier justification of expected utility by von Neumann and Morgenstern (1944) that had assumed known objective probabilities. Savage's axioms justify *subjective* assignments of probabilities. Even when accurate probabilities, such as the fifty-fifty put on the sides of a fair coin, are not available, decision makers conforming to Savage's axioms behave as if they form probabilities subjectively. Savage's axioms seem to undermine Knight's distinction between risk and uncertainty.

2.1. Savage and model misspecification. Savage's decision theory is both elegant and tractable. Furthermore, it provides a possible recipe for approaching concerns about model misspecification by putting a set of models on the table and averaging over them. For instance, think of a model as being a probability specification for the state of the world y tomorrow given the current state x and a decision or collection of decisions d : $f(y|x, d)$. If the conditional density f is unknown, then we can think about replacing f by a *family* of densities $g(y|x, d, \alpha)$ indexed by parameters α . By averaging over the array of candidate models using a prior (subjective) distribution, say π , we can form a 'hyper model' that we regard as being correctly specified. That is we can form:

$$f(y|x, d) = \int g(y|x, d, \alpha) d\pi(\alpha).$$

In this way, specifying the family of potential models and assigning a subjective probability distribution to them removes model misspecification.

Early examples of this so-called Bayesian approach to the analysis of policy-making in models with random coefficients are Friedman (1953) and Brainard (1967). The coefficient randomness can be viewed in terms of a subjective prior distribution. Recent developments in computational statistics have made this approach viable for a potentially rich class of candidate models.

This approach encapsulates specification concerns by formulating (1) a *set* of specific possible models, and (2) a prior distribution over those models. Below we raise questions about the extent to which these steps can really fully capture our concerns about model misspecification. As concerns (1), a hunch that a model is wrong might occur in a vague form that 'some other good fitting model actually governs the data' and that might not so readily translate into a well enumerated set of explicit and well formulated alternative models $g(y|x, d, \alpha)$. As concerns (2), even when we can specify a manageable set of well defined alternative models, we might

struggle to assign a unique prior $\pi(\alpha)$ to them. Hansen and Sargent (2007) address both of these concerns. They use a risk-sensitivity operator \mathbb{T}^1 as an alternative to (1) by taking each approximating model $g(y|x, d, \alpha)$, one for each α , and effectively surrounding each one with a cloud of models specified only in terms of how closely they approximate the conditional density $g(y|x, d, \alpha)$ statistically. Then they use a second risk-sensitivity operator \mathbb{T}^2 to surround a given prior $\pi(\alpha)$ with a set of priors that again are statistically close to the baseline π . We describe an application to a macroeconomic policy problem in section 5.4.

2.2. Savage and rational expectations. Rational expectations theory withdrew freedom from Savage's decision theory by imposing equality between agents' subjective probabilities and the probabilities emerging from the economic model containing those agents. Equating objective and subjective probability distributions removes all parameters that summarize agents' subjective distributions, and by doing so creates the powerful cross-equation restrictions characteristic of rational expectations empirical work.⁴ However, by insisting that subjective probabilities agree with objective ones, rational expectations make it much more difficult to dispose of Knight's distinction between risk and uncertainty by appealing to Savage's Bayesian interpretation of probabilities. Indeed, by equating objective and subjective probability distributions, the rational expectations hypothesis precludes a self-contained analysis of model misspecification. Because it abandons Savage's personal theory of probability, it can be argued that rational expectations indirectly increases the appeal of Knight's distinction between risk and uncertainty. Epstein and Wang (1994) argue that the Ellsberg paradox should make us rethink the foundation of rational expectations models.

2.3. The Ellsberg paradox. Ellsberg (1961) expressed doubts about the Savage approach by refining an example originally put forward by Knight. Consider two urns. In Urn A it is known that there are exactly ten red balls and ten black balls. In Urn B there are twenty balls, some red and some black. A ball from each urn is to be drawn at random. Free of charge, a person can choose one of the two urns and then place a bet on the color of the ball that is drawn. If he or she correctly guesses the color, the prize is 1 million dollars, while the prize is 0 dollars if the guess is incorrect. According to the Savage theory of decision-making, Urn B should be chosen even though the fraction of balls is not known. Probabilities can be formed subjectively, and a bet placed on the (subjectively) most likely ball color. If subjective probabilities are not fifty-fifty, a bet on Urn B will be strictly preferred to one on Urn A. If the subjective probabilities are precisely fifty-fifty then the decision maker will be indifferent. Ellsberg (1961) argued that a strict preference for Urn A is plausible because the probability of drawing a red or black ball is known in advance. He surveyed the preferences of an elite group of economists to lend support to this

⁴For example, see Sargent (1981).




FIGURE 1. The Ellsberg Urn

position.⁵ This example, called the *Ellsberg paradox*, challenges the appropriateness of the full array of Savage axioms.⁶

2.4. Multiple priors. Motivated in part by the Ellsberg (1961) paradox, Gilboa and Schmeidler (1989) provided a weaker set of axioms that included a notion of uncertainty aversion. Uncertainty aversion represents a preference for knowing probabilities over having to form them subjectively based on little information. Consider a choice between two gambles between which you are indifferent. Imagine forming a new bet that mixes the two original gambles with known probabilities. In contrast to von Neumann and Morgenstern (1944) and Savage (1954), Gilboa and Schmeidler (1989) did not require indifference to the mixture probability. Under *aversion to uncertainty*, mixing with known probabilities can only improve the welfare of the decision maker. Thus, Gilboa and Schmeidler required that the decision maker at least weakly prefer the mixture of gambles to either of the original gambles.

⁵Subsequent researchers have collected more evidence to substantiate this type of behavior. See Camerer (1999), especially table 3.2 on page 57, and also Harlevy (2007).

⁶In contrast to Ellsberg, Knight's second urn contained seventy-five red balls and twenty-five black balls (see Knight (1921), page 219). While Knight contrasted bets on the two urns made by different people, he conceded that if an action was to be taken involving the first urn, the decision maker would act under 'the supposition that the chances are equal.' He did not explore decisions involving comparisons of urns like that envisioned by Ellsberg.

The resulting generalized decision theory implies a *family* of priors and a decision maker who uses the worst case among this family to evaluate future prospects. Assigning a family of beliefs or probabilities instead of a unique prior belief renders Knight's distinction between risk and uncertainty operational. *After* a decision has been made, the family of priors underlying it can typically be reduced to a unique prior by averaging using subjective probabilities, from Gilboa and Schmeidler (1989). However, the prior that would be discovered by that procedure depends on the decision being considered and is an artifact of a decision-making process designed to make a conservative assessment. In the case of the Knight-Ellsberg urn example, a range of priors is assigned to red balls, say .45 to .55, and similarly to black balls in Urn B. The conservative assignment of .45 to red balls when evaluating a red ball bet and .45 to black balls when making a black ball bet implies a preference for Urn A. A bet on either ball color from Urn A has a .5 probability of success.

A product of the Gilboa-Schmeidler axioms is a decision theory that can be formalized as a two-player game. For every action of one maximizing player, a second minimizing player selects associated beliefs. The second player chooses those beliefs in a way that balances the first player's wish to make good forecasts against his doubts about model specification.⁷

Just as the Savage axioms do not tell a model-builder how to specify the subjective beliefs of decision makers for a given application, the Gilboa-Schmeidler axioms do not tell a model-builder the *family* of potential beliefs. The axioms only clarify the sense in which rational decision making may require multiple priors along with a fictitious second agent who selects beliefs in a pessimistic fashion. Restrictions on beliefs must come from outside.⁸

2.5. Ellsberg and Friedman. The Knight-Ellsberg Urn example might look far removed from the dynamic models used in macroeconomics. But a fascinating chapter in the history of macroeconomics centers on Milton Friedman's ambivalence about expected utility theory. Although Friedman embraced the expected utility theory of von Neumann and Morgenstern (1944) in some work (Friedman and Savage (1948)), he chose not to use it⁹ when discussing the conduct of monetary policy. Instead, Friedman (1959) emphasized that model misspecification is a decisive consideration for monetary and fiscal policy. Discussing the relation between money and prices, Friedman concluded that:

⁷The theory of zero-sum games gives a natural way to make a concern about robustness algorithmic. Zero-sum games were used in this way in both statistical decision theory and robust control theory long before Gilboa and Schmeidler supplied their axiomatic justification. See Blackwell and Girshick (1954), Ferguson (1967), and Jacobson (1973).

⁸That, of course, was why restriction-hungry macroeconomists and econometricians seized on the ideas of Muth (1961) in the first place.

⁹Unlike Lucas (1976) and Sargent and Wallace (1975).

If the link between the stock of money and the price level were direct and rigid, or if indirect and variable, fully understood, this would be a distinction without a difference; the control of one would imply the control of the other; But the link is not direct and rigid, nor is it fully understood. While the stock of money is systematically related to the price level *on the average*, there is much variation in the relation over short periods of time Even the variability in the relation between money and prices would not be decisive if the link, though variable, were synchronous so that current changes in the stock of money had their full effect on economic conditions and on the price level instantaneously or with only a short lag. ... In fact, however, there is much evidence that monetary changes have their effect only after a considerable lag and over a long period and that lag is rather variable.

Friedman thought that misspecification of the dynamic link between money and prices should concern proponents of activist policies. Despite Friedman and Savage (1948), his treatise on monetary policy (Friedman (1959)) did not advocate forming prior beliefs over alternative specifications of the dynamic models in response to this concern about model misspecification.¹⁰ His argument reveals a preference not to use Savage's decision theory for the practical purpose of designing monetary policy.

3. FORMALIZING A TASTE FOR ROBUSTNESS

The multiple priors formulation provides a way to think about model misspecification. Like Epstein and Wang (1994) and Friedman (1959), we are specifically interested in decision-making in dynamic environments. We draw our inspiration from a line of research in control theory. Robust control theorists challenged and reconstructed earlier versions of control theory because it had ignored model-approximation error in designing policy rules. They suspected that their models had misspecified the dynamic responses of target variables to controls. To confront that concern, they added a specification error process to their models, and sought decision rules that would work well across a set of such error processes. That led them to a two-player game and a conservative-case analysis much in the spirit of Gilboa and Schmeidler (1989). In this section, we describe the modifications of modern control theory made by the robust control theorists. While we feature linear/quadratic Gaussian control, many of the results that we discuss have direct extensions to more general decision environments. For instance, Hansen, Sargent, Turmuhambetova, and Williams (2006) consider robust decision problems in Markov diffusion environments.

¹⁰However, Friedman (1953) conducts an explicitly stochastic analysis of macroeconomic policy and introduces elements of the analysis of Brainard (1967).

3.1. Control with a correct model. First, we briefly review standard control theory, which does not admit misspecified dynamics. For pedagogical simplicity, consider the following state evolution and target equations for a decision maker:

$$(1) \quad x_{t+1} = Ax_t + Bu_t + Cw_{t+1}$$

$$(2) \quad z_t = Hx_t + Ju_t$$

where x_t is a state vector, u_t is a control vector, and z_t is a target vector, all at date t . In addition, suppose that $\{w_{t+1}\}$ is a sequence of vectors of independent and identically and normally distributed *shocks* with mean zero and covariance matrix given by I . The target vector is used to define preferences via:

$$(3) \quad -\frac{1}{2} \sum_{t=0}^{\infty} \beta^t E z_t' z_t$$

where $0 < \beta < 1$ is a discount factor and E is the mathematical expectation operator. The aim of the decision maker is to maximize this objective function by choice of control law $u_t = -Fx_t$. The linear form of this decision rule for u_t is not a restriction but is an implication of optimality.

The explicit, stochastic, recursive structure makes it tractable to solve the control problem via dynamic programming:

Problem 1. (*Recursive Control*)

Dynamic programming reduces this infinite-horizon control problem to the following fixed-point problem in the matrix Ω in the following functional equation:

$$(4) \quad -\frac{1}{2}x'\Omega x - \omega = \max_u \left\{ -\frac{1}{2}z'z - \frac{\beta}{2}Ex^*\Omega x^* - \beta\omega \right\}$$

subject to

$$x^* = Ax + Bu + Cw^*$$

where w^* has mean zero and covariance matrix I .¹¹ Here * superscripts denote next-period values.

The solution of the ordinary linear quadratic optimization problem has a special property called *certainty equivalence* that asserts that the decision rule F is independent of the volatility matrix C . We state this formally in the following claim:

Claim 2. (*Certainty Equivalence Principle*)

For the linear-quadratic control problem, the matrix Ω and the optimal control law F do not depend on the volatility matrix C . Thus, the optimal control law does not depend on the matrix C .

¹¹There are considerably more computationally efficient solution methods for this problem. See Anderson, Hansen, McGrattan, and Sargent (1996) for a survey.

The certainty equivalence principle comes from the quadratic nature of the objective, the linear form of the transition law, and the specification that the *shock* w^* is independent of the current state x . Robust control theorists challenge this solution because of their experience that it is vulnerable to model misspecification. Seeking control rules that will do a good job for a *class* of models induces them to focus on alternative possible shock processes.

Can a temporally independent shock process w_{t+1} represent the kinds of misspecification decision makers fear? Control theorists think not, because they fear misspecified *dynamics*, i.e., misspecifications that affect the impulse response functions of target variables to shocks and controls. For this reason, they formulate misspecification in terms of shock processes that can feed back on the state variables, something that i.i.d. shocks cannot do. As we shall see, allowing the shock to feed back on current and past states will modify the certainty equivalence property.

3.2. Model misspecification. To capture misspecification in the dynamic system, suppose that the i.i.d. shock sequence is replaced by *unstructured* model specification errors. We temporarily replace the stochastic shock process $\{w_{t+1}\}$ with a deterministic sequence $\{v_t\}$ of model approximation errors of limited magnitude. As in Gilboa and Schmeidler (1989), a two-person zero-sum game can be used to represent a preference for decisions that are robust with respect to v . We have temporarily suppressed randomness, so now the game is dynamic and deterministic.¹² As we know from the dynamic programming formulation of the single-agent decision problem, it is easier to think of this problem recursively. A value function conveniently encodes the impact of current decisions on future outcomes.

Game 3. (*Robust Control*)

To represent a preference for robustness, we replace the single-agent maximization problem (4) by the two-person dynamic game:

$$(5) \quad -\frac{1}{2}x'\Omega x = \max_u \min_v -\frac{1}{2}z'z + \frac{\theta}{2}v'v - \frac{\beta}{2}x^{*'}\Omega x^*$$

subject to

$$x^* = Ax + Bu + Cv$$

where $\theta > 0$ is a parameter measuring a preference for robustness. Again we have formulated this as a fixed-point problem in the value function: $V(x) = -\frac{1}{2}x'\Omega x - \omega$.

Notice that a malevolent agent has entered the analysis. This agent, or alter ego, aims to *minimize* the objective, but in doing so is penalized by a term $\frac{\theta}{2}v'v$ that is added to the objective function. Thus, the theory of dynamic games can be applied to study robust decision-making, a point emphasized by Basar and Bernhard (1995).

¹²See appendix A for an equivalent but more basic stochastic formulation of the following robust control problem.

The fictitious second agent puts context-specific *pessimism* into the control law. Pessimism is context specific and endogenous because it depends on the details of the original decision problem, including the one-period return function and the state evolution equation. The robustness parameter or multiplier θ restrains the magnitude of the pessimistic distortion. Large values of θ keep the degree of pessimism (the magnitude of v) small. By making θ arbitrarily large, we approximate the certainty-equivalent solution to the single-agent decision problem.

3.3. Types of misspecifications captured. In formulation (5), the solution makes v a function of x and u a function of x alone. Associated with the solution to the two-player game is a worst-case choice of v . The dependence of the “worst-case” model shock v on the control u and the state x is used to promote robustness. This worst case corresponds to a particular (A^\dagger, B^\dagger) that is a device to acquire a robust rule. If we substitute the value-function fixed point into the right side of (5) and solve the inner minimization problem, we obtain the following formula for the worst-case error:

$$(6) \quad v^\dagger = (\theta I - \beta C' \Omega C)^{-1} C' \Omega (Ax + Bu).$$

Notice that this v^* depends on both the current period control vector u and state vector x . Thus, the misspecified model used to promote robustness has:

$$\begin{aligned} A^\dagger &= A + C(\theta I - \beta C' \Omega C)^{-1} C' \Omega A \\ B^\dagger &= B + C(\theta I - \beta C' \Omega C)^{-1} C' \Omega B. \end{aligned}$$

Notice that the resulting distorted model is context specific and depends on the matrices A, B, C , the matrix Ω used to represent the value function, and the robustness parameter θ .

The matrix Ω is typically positive semidefinite, which allows us to exchange the maximization and minimization operations:

$$(7) \quad -\frac{1}{2} x' \Omega x = \min_v \max_u -\frac{1}{2} z' z + \frac{\theta}{2} v' v - \frac{\beta}{2} x^{*'} \Omega x^*$$

We obtain the same value function even though now u is chosen as a function of v and x while v depends only on x . For this solution:

$$u^\dagger = -(J' J + B' \Omega B)^{-1} J' [Hx + \Omega (Ax + Cv)]$$

The equilibrium v that emerges in this alternative formulation gives an alternative dynamic evolution equation for the state vector x . The robust control u is a best response to this alternative evolution equation (given Ω). In particular, abusing notation, the alternative evolution is:

$$x^* = Ax + Cv(x) + Bu$$

The equilibrium outcomes from zero-sum games (5) and (7) in which both v and u are represented as functions of x alone coincide.

This construction of a worst-case model by exchanging orders of minimization and maximization may sometimes be hard to interpret as a plausible alternative model. Moreover, the construction depends on the matrix Ω from the recursive solution to the robust control problem and hence includes a contribution from the penalty term. As an illustration of this problem, suppose that one of the components of the state vector is exogenous, by which we mean a state vector that cannot be influenced by the choice of the control vector. But under the alternative model this component may fail to be exogenous. The alternative model formed from the worst-case shock $v(x)$ as described above may thus include a form of endogeneity that is hard to interpret. Hansen and Sargent (2008b) describe ways to circumvent this annoying apparent endogeneity by an appropriate application of the macroeconomist's 'Big K, little k' trick.¹³

What legitimizes the exchange of minimization and maximization in the recursive formulation is something referred to as a Bellman-Isaacs condition. When this condition is satisfied, we can exchange orders in the date zero problem. This turns out to give us an alternative construction of a worst-case model that can avoid any unintended endogeneity of the worst-case model. In addition, the Bellman-Isaacs condition is central in justifying the use of recursive methods for solving date-zero robust control problems. See the discussions in Fleming and Souganidis (1989), Hansen, Sargent, Turmuhambetova, and Williams (2006), and Hansen and Sargent (2008b).

What was originally the volatility exposure matrix C now also becomes an impact matrix for misspecification. It contributes to the solution of the robust control problem, while for the ordinary control problem, it did not by virtue of certainty equivalence. We summarize the dependence of F on C in the following, which is fruitfully compared and contrasted with claim 2:

Claim 4. (*Breaking Certainty Equivalence*)

For $\theta < +\infty$, the robust control $u = -Fx$ that solves game (3) depends on the volatility matrix C .

We shall remark below how the breaking down of certainty equivalence is attributable to a kind of precautionary motive emanating from fear of model misspecification. While the certainty equivalent benchmark is special, it points to a force prevalent in more general settings. Thus, in settings where the presence of random shocks *does* have an impact on decision rules in the absence of a concern about misspecification, introducing such concerns typically leads to an enhanced precautionary motive.

3.4. Gilboa and Schmeidler again. To relate formulation (3) to that of Gilboa and Schmeidler (1989), we look at a specification in which we alter the distribution of the shock vector. The idea is to change the conditional distribution of the shock

¹³See Ljungqvist and Sargent (2004), p. 384.

vector from a multivariate standard normal that is independent of the current state vector by multiplying this baseline density by a likelihood ratio (relative to the standardized multivariate normal). This likelihood ratio can depend on current and past information in a general fashion so that general forms of misspecified dynamics can be entertained when solving versions of a two-player zero-sum game in which the minimizing player chooses the distorting density. This more general formulation allows misspecifications that include neglected nonlinearities, higher-order dynamics, and an incorrect shock distribution. As a consequence, this formulation of robustness is called *unstructured*.¹⁴

For the linear-quadratic-Gaussian problem, it suffices to consider only changes in the conditional mean and the conditional covariance matrix of the shocks. See Appendix A for details. The worst-case covariance matrix is independent of the current state but the worst-case mean will depend on the current state. This conclusion extends to continuous-time decision problems that are not linear-quadratic provided that the underlying shocks can be modeled as diffusion processes. It suffices to explore misspecifications that append state dependent drifts to the underlying Brownian motions. See Hansen, Sargent, Turmuhambetova, and Williams (2006) for a discussion. The quadratic penalty $\frac{1}{2}v'v$ becomes a measure of what is called *conditional relative entropy* in the applied mathematic literature. It is a discrepancy measure between an alternative conditional density and, for example, the normal density in a baseline model. Instead of restraining the alternative densities to reside in some prespecified set, for convenience we penalize their magnitude directly in the objective function. As discussed in Hansen, Sargent, and Tallarini (1999), Hansen, Sargent, Turmuhambetova, and Williams (2006) and Hansen and Sargent (2008b), we can think of the robustness parameter θ as a Lagrange multiplier on a time 0 constraint on discounted relative entropy.¹⁵

4. CALIBRATING A TASTE FOR ROBUSTNESS

Our model of a robust decision maker is formalized as a two-person zero-sum dynamic game. The minimizing player, if left unconstrained, can inflict serious damage and substantially alter the decision rules. It is easy to construct examples in which the induced conservative behavior is so cautious that it makes the robust decision rule look silly. Such examples can be used to promote skepticism about the use of minimization over models rather than the averaging advocated in Bayesian decision theory.

¹⁴See Onatski and Stock (1999) for an example of robust decision analysis with structured uncertainty.

¹⁵See Hansen and Sargent (2001), Hansen, Sargent, Turmuhambetova, and Williams (2006), and Hansen and Sargent (2008b), chapter 7, for discussions of ‘multiplier’ preferences defined in terms of θ and ‘constraint preferences’ that are special cases of preferences supported by the axioms of Gilboa and Schmeidler (1989).

Whether the formulation in terms of the two-person zero-sum game looks silly or plausible depends on how the choice set open to the fictitious minimizing player is disciplined. While an undisciplined malevolent player can wreak havoc, a tightly constrained one cannot. Thus, the interesting question is whether it is reasonable as either a positive or normative model of decision-making to make conservative adjustments induced by ambiguity over model specification, and if so, how big these adjustments should be. Some support for making conservative adjustments appears in experimental evidence (see Camerer (1995) for a discussion) and other support comes from the axiomatic treatment of Gilboa and Schmeidler (1989). Neither of these sources answer the quantitative question of how *large* the adjustment should be in applied work in economic dynamics. Here we think that the theory of statistical discrimination can help.

We have parameterized a taste for robustness in terms of a single free parameter, θ , or else implicitly in terms of the associated discounted entropy η_0 . Let M_t denote the date t likelihood ratio of an alternative model vis a vis the original “approximating” model. Then $\{M_t : t = 0, 1, \dots\}$ is a martingale under the original probability law, and we normalize $M_0 = 1$. The date zero measure of relative entropy is

$$E(M_t \log M_t | \mathcal{F}_0),$$

which is the expected log-likelihood ratio under the alternative probability measure, where \mathcal{F}_0 is the information set at time 0. For infinite-horizon problems, we find it convenient to form a geometric average using the subjective discount factor $\beta \in (0, 1)$ to construct the geometric weights,

$$(8) \quad (1 - \beta) \sum_{j=0}^{\infty} \beta^j E(M_j \log M_j | \mathcal{F}_0) \leq \eta_0.$$

By a simple summation-by-parts argument,

$$(9) \quad (1 - \beta) \sum_{j=0}^{\infty} \beta^j E(M_j \log M_j | \mathcal{F}_0) = \sum_{j=0}^{\infty} \beta^j E(M_j \log M_j - \log M_{j-1} | \mathcal{F}_0).$$

For computational purposes it is useful to use a penalization approach and to solve the decision problems for alternative choices of θ . Associated with each θ , we can find a corresponding value of η_0 . This seemingly innocuous computational simplification has subtle implications for the specification of preferences. In defining preferences, it matters if you hold fixed θ (here you get the so-called multiplier preferences) or hold fixed η_0 (and here you get the so-called constraint preferences.) See Hansen, Sargent, Turmuhambetova, and Williams (2006) and Hansen and Sargent (2008b) for discussions. Even when we adopt the multiplier interpretation of preferences, it is revealing to compute the implied η_0 's as suggested by Petersen, James, and Dupuis (2000).

For the purposes of calibration we want to know which values of the parameter θ correspond to *reasonable* preferences for robustness. To think about this issue, we start by recalling that the rational expectations notion of equilibrium makes the model that economic agents use in their decision-making be the same model that generates the observed data. A defense of the rational expectations equilibrium concept is that discrepancies between models should have been detected from sufficient historical data and then eliminated. In this section, we use a closely related idea to think about reasonable preferences for robustness. Given historical observations on the state vector, we use a Bayesian model detection theory originally due to Chernoff (1952). This theory describes how to discriminate between two models as more data become available. We use statistical detection to limit the preference for robustness. The decision maker should have noticed easily detected forms of model misspecification from past time series data and eliminated them. We propose restricting θ to admit only alternative models that are difficult to distinguish statistically from the approximating model. We do this rather than study a considerably more complicated learning and control problem. We will discuss relationships between robustness and learning in section 5.

4.1. State evolution. Given a time series of observations on the state vector x_t , suppose that we want to determine the evolution equation for the state vector. Let $u = -F^\dagger x$ denote the solution to the robust control problem. One possible description of the time series is

$$(10) \quad x_{t+1} = (A - BF^\dagger)x_t + Cw_{t+1}$$

where $\{w_{t+1}\}$ is a sequence of i.i.d. normalized Gaussian vectors. In this case, concerns about model misspecification are just in the head of the decision maker: the original model is actually correctly specified. Here the approximating model actually generates the data.

An worst-case evolution equation is the one associated with the solution to the two-player zero-sum game. This changes the distribution of w_{t+1} by appending a conditional mean as in (6)

$$v^\dagger = -K^\dagger x$$

where

$$K^\dagger = \frac{1}{\theta} \left(I - \frac{\beta}{\theta} C' \Omega^* C \right)^{-1} C' \Omega^* (A - BF^r).$$

and altering the covariance matrix CC' . The alternative evolution remains Markov and can be written as:

$$(11) \quad x_{t+1} = (A - BF^\dagger - CK^\dagger)x_t + Cw_{t+1}^\dagger.$$

where

$$w_{t+1} = -K^\dagger x_t + w_{t+1}^\dagger$$

and w_{t+1}^\dagger is normally distributed with mean zero, but a covariance matrix that typically exceeds the identity matrix. This evolution takes the constrained worst-case model as the actual law of motion of the state vector, evaluated under the robust decision rule and the worst-case shock process that the decision maker plans against.¹⁶ Since the choice of v by the minimizing player is not meant to be a prediction, only a conservative adjustment, this evolution equation is not the decision maker's guess about the most likely model. The decision maker considers more general changes in the distribution for the shock vector w_{t+1} , but the implied relative entropy (9) is no larger than that for the model just described. The actual misspecification could take on a more complicated form than the solution to the two-player zero-sum game. Nevertheless, the two evolution equations (10) and (11) provide a convenient laboratory for calibrating plausible preferences for robustness.

4.2. Classical model detection. The log-likelihood ratio is used for statistical model selection. For simplicity, consider pairwise comparisons between models. Let one be the basic approximating model captured by (A, B, C) and a multivariate standard normal shock process $\{w_{t+1}\}$. Suppose another is indexed by $\{v_t\}$ where v_t is the conditional mean of w_{t+1} . The underlying randomness masks the model misspecification and allows us to form likelihood functions as a device for studying how informative data are in revealing which model generates the data.¹⁷

Imagine that we observe the state vector for a finite number T of time periods. Thus, we have x_1, x_2, \dots, x_T . Form the log likelihood ratio between these two models. Since the $\{w_{t+1}\}$ sequence is independent and identically normally distributed, the date t contribution to the log likelihood ratio is

$$w_{t+1} \cdot \hat{v}_t - \frac{1}{2} \hat{v}_t \cdot \hat{v}_t$$

where \hat{v}_t is the modeled version of v_t . For instance, we might have that $\hat{v}_t = f(x_t, x_{t-1}, \dots, x_{t-k})$. When the approximating model is correct, $v_t = 0$ and the predictable contribution to the (log) likelihood function is negative: $-\frac{1}{2} \hat{v}_t \cdot \hat{v}_t$. When the alternative \hat{v}_t model is correct, the predictable contribution is $\frac{1}{2} \hat{v}_t \cdot \hat{v}_t$. Thus, the term $\frac{1}{2} \hat{v}_t \cdot \hat{v}_t$ is the average (conditioned on current information) time t contribution to a log-likelihood ratio. When this term is large, model discrimination is easy, but it is difficult when this term is small. This motivates our use of the quadratic form $\frac{1}{2} \hat{v}_t \cdot \hat{v}_t$ as a statistical measure of model misspecification. Of course, the \hat{v}_t 's depend

¹⁶It is the decision rule from the Markov perfect equilibrium of the dynamic game.

¹⁷Here, for pedagogical convenience we explore only a special stochastic departure from the approximating model. As emphasized by Anderson, Hansen, and Sargent (2003), statistical detection theory leads us to consider only model departures that are absolutely continuous with respect to the benchmark or approximating model. The departures considered here are the discrete-time counterparts to the departures admitted by absolute continuity when the state vector evolves according to a possibly nonlinear diffusion model.

on the state x_t , so that to simulate them requires simulating a particular law of motion (11).

Use of $\frac{1}{2}\hat{v}_t \cdot \hat{v}_t$ as a measure of discrepancy is based implicitly on a classical notion of statistical discrimination. Classical statistical practice typically holds fixed the type I error of rejecting a given *null* model when the null model is true. For instance, the null model might be the benchmark \hat{v}_t model. As we increase the amount of available data, the type II error of accepting the null model when it is false decays to zero as the sample size increases, typically at an exponential rate. The likelihood-based measure of model discrimination gives a lower bound on the rate (per unit observation) at which the type II error probability decays to zero.

4.3. Bayesian model detection. Chernoff (1952) studied a Bayesian model discrimination problem. Suppose we average over both the type I and II errors by assigning prior probabilities of say one half to each model. Now additional information at date t allows one to improve model discrimination by shrinking both type I and type II errors. This gives rise to a discrimination rate (the deterioration of log probabilities of making a classification error per unit time) equal to $\frac{1}{8}\hat{v}_t \cdot \hat{v}_t$ for the Gaussian model with only differences in means, although Chernoff entropy is defined much more generally. This rate is known as Chernoff entropy. When the Chernoff entropy is small, models are hard to tell apart statistically. When Chernoff entropy is large, statistical detection is easy. The scaling by $\frac{1}{8}$ instead of $\frac{1}{2}$ reflects the trade-off between type I and type II errors. Type I errors are no longer held constant. Notice that the penalty term that we added to the control problem to enforce robustness is a scaled version of Chernoff entropy, provided that the model misspecification is appropriately disguised by Gaussian randomness. Thus, when thinking about statistical detection, it is imperative that we include some actual randomness, which though absent in many formulations of robust control theory, is present in virtually all macroeconomic applications.

In a model generating data that are independent and identically distributed, we can accumulate the Chernoff entropies over the observation indices to form a detection error probability bound for finite samples. In dynamic contexts, more is required than just this accumulation, but it is still true that Chernoff entropy acts as a *short-term* discount rate in the construction of the probability bound.¹⁸

We believe that the model detection problem confronted by a decision maker is actually more complicated than the pairwise statistical discrimination problem we just described. A decision maker will most likely be concerned about a wide array of more complicated models, many of which may be more difficult to formulate and solve than the ones considered here. Nevertheless, this highly stylized framework for statistical discrimination gives one way to think about a plausible preference for robustness. For any given θ , we can compute the implied worst-case process $\{v_t^\dagger\}$

¹⁸See Anderson, Hansen, and Sargent (2003).

and consider only those values of θ for which the $\{v_t^\dagger\}$ model is hard to distinguish from the $v_t = 0$ model. From a statistical standpoint, it is more convenient to think about the magnitude of the v_t^\dagger 's than of the θ 's that underlie them. This suggests solving robust control problems for a set of θ 's and exploring the resulting v_t^\dagger 's. Indeed, Anderson, Hansen, and Sargent (2003) establish a close connection between $v_t^\dagger \cdot v_t^\dagger$ and (a bound on) a detection error probability.

4.3.1. *Detection probabilities: an example.* Here is how we construct detection error probabilities in practice. Consider two alternative models with equal prior probabilities. Model A is the approximating model and model B is the worst-case model associated with an alternative distribution for the shock process for a particular positive θ . Consider a fixed sample of T observations on x_t . Let L_i be the likelihood of that sample for model i for $i = A, B$. Define the likelihood ratio

$$\ell = \log L_A - \log L_B$$

We can draw a sample value of this log-likelihood ratio by generating a simulation of length T for x_t under model i . The Bayesian detection error probability averages probabilities of two kinds of errors. First, assume that model A generates the data and calculate

$$p_A = \text{Prob}(\text{error}|A) = \text{freq}(\ell \leq 0|A).$$

Next, assume that model B generates the data and calculate

$$p_B = \text{Prob}(\text{error}|B) = \text{freq}(\ell \geq 0|B).$$

Since the prior equally weights the two models, the probability of a detection error is

$$p(\theta) = \frac{1}{2}(p_A + p_B).$$

Our idea is to set $p(\theta)$ at a plausible value, then to invert $p(\theta)$ to find a plausible value for the preference-for-robustness parameter θ . We can approximate the values of p_A, p_B composing $p(\theta)$ by simulating a large number N of realizations of samples of x_t of length T . In the example below, we simulated 20,000 samples. See Hansen, Sargent, and Wang (2002) for more details about computing detection error probabilities.

We now illustrate the use of detection error probabilities to discipline the choice of θ in the context of the simple dynamic model that Ball (1999) designed to study alternative rules by which a monetary policy authority might set an interest rate.¹⁹

¹⁹See Sargent (1999a) for further discussion of Ball's model from the perspective of robust decision theory. See Hansen and Sargent (2008b) (chapter 16) for how to treat robustness in 'forward-looking' models.

Ball's is a 'backward-looking' macro model with the structure

$$(12) \quad y_t = -\beta r_{t-1} - \delta e_{t-1} + \epsilon_t$$

$$(13) \quad \pi_t = \pi_{t-1} + \alpha y_{t-1} - \gamma(e_{t-1} - e_{t-2}) + \eta_t$$

$$(14) \quad e_t = \theta r_t + \nu_t,$$

where y is the logarithm of real output, r is the real interest rate, e is the logarithm of the real exchange rate, π is the inflation rate, and ϵ, η, ν are serially uncorrelated and mutually orthogonal disturbances. As an objective, Ball assumed that a monetary authority wants to maximize

$$-E(\pi_t^2 + y_t^2).$$

The monetary authority sets the interest rate r_t as a function of the current state, which Ball shows can be reduced to y_t, e_t .

Ball motivates (12) as an open-economy IS curve and (13) as an open-economy Phillips curve; he uses (14) to capture effects of the interest rate on the exchange rate. Ball set the parameters $\gamma, \theta, \beta, \delta$ to the values .2, 2, .6, .2. Following Ball, we set the innovation shock standard deviations equal to 1, 1, $\sqrt{2}$, respectively.

To discipline the choice of the parameter expressing a preference for robustness, we calculated the detection error probabilities for distinguishing Ball's model from the worst-case models associated with various values of $\sigma \equiv -\theta^{-1}$. We calculated these taking Ball's parameter values as the approximating model and assuming that $T = 142$ observations are available, which corresponds to 35.5 years of annual data for Ball's quarterly model. Figure 2 shows these detection error probabilities $p(\sigma)$ as a function of σ . Notice that the detection error probability is .5 for $\sigma = 0$, as it should be, because then the approximating model and the worst-case model are identical. The detection error probability falls to .1 for $\sigma \approx -.085$. If we think that a reasonable preference for robustness is to design rules that work well for alternative models whose detection error probabilities are .1 or greater, then $\sigma = -.085$ is a reasonable choice of this parameter. Later, we will compute a robust decision rule for Ball's model with $\sigma = -.085$ and compare its performance to the $\sigma = 0$ rule that expresses no preference for robustness.

4.3.2. Reservations and extensions. Our formulation treats misspecification of all of the state-evolution equations symmetrically and admits all misspecification that can be disguised by the shock vector w_{t+1} . Our hypothetical statistical discrimination problem assumes historical data sets of a common length on the entire state vector process. We might instead imagine that there are differing amounts of confidence in state equations not captured by the perturbation Cv_t and quadratic penalty $\theta v_t \cdot v_t$. For instance, to imitate aspects of Ellsberg's two urns we might imagine that misspecification is constrained to be of the form $C \begin{bmatrix} v_t^1 \\ 0 \end{bmatrix}$ with corresponding penalty $\theta v_t^1 \cdot v_t^1$. The rationale for the restricted perturbation would be that there is more confidence in some aspects of the model than others. More generally, multiple




FIGURE 2. Detection error probability (coordinate axis) as a function of $\sigma = -\theta^{-1}$ for Ball's model.

penalty terms could be included with different weighting. A cost of this generalization is a greater burden on the calibrator. More penalty parameters would need to be selected to model a robust decision maker.

The preceding use of the theory of statistical discrimination conceivably helps to excuse a decision not to model active learning about model misspecification. But sometimes that excuse might not be convincing. For that reason, we next explore ways of incorporating learning.

5. LEARNING

The robust control model outlined above sees decisions being made via a two-stage process:

- 1. There is an initial learning-model-specification period during which data are studied and an approximating model is specified. This process is taken for granted and not analyzed. However, afterwards, learning ceases, though doubts surround the model specification.
- 2. Given the approximating model, a single fixed decision rule is chosen and used forever. Though the decision rule is designed to guard against model misspecification, no attempt is made to use the data to narrow the model ambiguity during the control period.

The defense for this two-stage process is that somehow the first stage discovers an approximating model and a set of surrounding models that are difficult to distinguish

from it with the data available in stage 1 and that are likely to be available only after a long time has passed in stage 2.

This section considers approaches to model ambiguity coming from literatures on adaptation and that do not temporally separate learning from control as in the two step process just described. Instead, they assume continuous learning about the model and continuous adjustment of decision rules.

5.1. Bayesian models. For a low-dimensional specification of model uncertainty, an explicit Bayesian formulation might be an attractive alternative to our robust formulation. We could think of matrices A and B in the state evolution (1) as being random and specify a prior distribution for this randomness. One possibility is that there is only some initial randomness, to represent the situation that A and B are unknown but fixed in time. In this case, observations of the state would convey information about the realized A and B . Given that the controller does not observe A and B , and must make inference about these matrices as time evolves, this problem is not easy to solve. Nevertheless, numerical methods may be employed to approximate solutions. For example, see Wieland (1996) and Cogley, Colacito, and Sargent (2007).

We shall use a setting of Cogley, Colacito, and Sargent (2007) first to illustrate purely Bayesian procedures for approaching model uncertainty, then to show how to adapt these to put robustness into decision rules. A decision maker wants to maximize the following function of states s_t and controls v_t :

$$(15) \quad E_0 \sum_{t=0}^{\infty} \beta^t r(s_t, v_t).$$

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

$$(16) \quad s_{t+1} = g(s_t, v_t, z_t, \epsilon_{t+1}),$$

$$(17) \quad z_{t+1} = z_t,$$

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 (16)-(17) 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 $\text{Prob}(z = 1) = \pi_0$. Given history $s^t = [s_t, s_{t-1}, \dots, s_0]$, the decision maker recursively computes $\pi_t = \text{Prob}(z = 1 | s^t)$ by applying Bayes' law:

$$(18) \quad \pi_{t+1} = B(\pi_t, g(s_t, v_t, z_t, \epsilon_{t+1})).$$

For example, Cogley, Colacito, Hansen, and Sargent (2008) take one of the submodels to be a Keynesian model of a Phillips curve while the other is a new classical model. The decision maker must decide while he learns.

Because he does not know z_t , the policy maker's prior probability π_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

$$(19) \quad V(s, \pi) = \max_v \left\{ r(s, v) + E_z \left[E_{s^*, \pi^*} (\beta V(s^*, \pi^*) | s, v, \pi, z) | s, v, \pi \right] \right\},$$

subject to

$$(20) \quad s^* = g(s, v, z, \epsilon^*),$$

$$(21) \quad \pi^* = B(\pi, g(s, v, z, \epsilon^*)).$$

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

5.2. Experimentation with specification doubts. Bellman equation (19) expresses the motivation that a decision maker has to experiment, i.e., to take into account how his decision affects future values of the component of the state π^* . We describe how Hansen and Sargent (2007) and Cogley, Colacito, Hansen, and Sargent (2008) adjust Bayesian learning and decision making to account for fears of model misspecification. Bellman equation (19) invites us to consider two types of misspecification of the stochastic structure: misspecification of the distribution of (s^*, π^*) conditional on (s, v, π, z) , and misspecification of the probability π over submodels z . Following Hansen and Sargent (2007), we introduce two risk-sensitivity operators that can help a 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.

5.3. Two risk-sensitivity operators.

5.3.1. \mathbb{T}^1 operator. The risk-sensitivity operator \mathbb{T}^1 helps the decision maker guard against misspecification of a submodel.²⁰ Let $W(s^*, \pi^*)$ be a measurable function of (s^*, π^*) . In our application, W will be a continuation value function. Instead of taking conditional expectations of W , Cogley, Colacito, Hansen, and Sargent (2008) and Hansen and Sargent (2007) apply the operator:

$$(22) \quad \mathbb{T}^1(W(s^*, \pi^*))(s, \pi, v, z; \theta_1) = -\theta_1 \log E_{s^*, \pi^*} \exp\left(\frac{-W(s^*, \pi^*)}{\theta_1}\right) \Big| (s, \pi, v, z)$$

²⁰See appendix A for more discussion on how to derive and interpret the risk-sensitivity operator \mathbb{T} .

where E_{s^*, π^*} denotes a mathematical expectation with respect to the conditional distribution of s^*, π^* . This operator yields the indirect utility function for a problem in which the minimizing agent chooses a worst-case distortion to the conditional distribution for (s^*, π^*) 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^*, π^*) . The worst-case distortion is proportional to

$$(23) \quad \exp\left(\frac{-W(s^*, \pi^*)}{\theta_1}\right),$$

where the factor of proportionality is chosen to make this nonnegative 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. A likelihood ratio proportional to (23) pessimistically twists the conditional density of (s^*, π^*) by upweighting outcomes that have lower continuation values.

5.3.2. 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, π, v, z) in a way that guards against misspecification of his prior π :

$$(24) \quad T^2(\widetilde{W}(s, \pi, v, z))(s, \pi, v; \theta_2) = -\theta_2 \log E_z \exp\left(\frac{-\widetilde{W}(s, \pi, v, z)}{\theta_2}\right) \Big|_{(s, \pi, v)}$$

This operator yields the indirect utility function for a problem in which the malevolent agent chooses a distortion to the Bayesian prior π in order to minimize the expected value of a function $\widetilde{W}(s, \pi, v, z)$ 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

$$(25) \quad \exp\left(\frac{-\widetilde{W}(s, \pi, v, z)}{\theta_2}\right),$$

where the factor of proportionality is chosen to make this nonnegative random variable have mean one. The worst-case density distorts the Bayesian prior by putting higher probability on outcomes with lower continuation values.

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.²¹ 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.

Application of this second risk-sensitivity operator provides a response to Levin and Williams (2003) and Onatski and Williams (2003). Levin and Williams (2003) explore multiple benchmark models. Uncertainty across such models can be expressed conveniently by the \mathbb{T}^2 operator and a concern for this uncertainty is implemented by making robust adjustments to model averages based on historical data.²² As is the aim of Onatski and Williams (2003), the \mathbb{T}^2 operator can be used to explore the consequences of unknown parameters as a form of “structured” uncertainty that is difficult to address via application of the \mathbb{T}^1 operator.²³ Finally application of the \mathbb{T}^2 operation gives a way to provide a benchmark to which one can compare the Taylor rule and other simple monetary policy rules.²⁴

5.4. A Bellman equation for inducing robust decision rules. Following Hansen and Sargent (2007), Cogley, Colacito, Hansen, and Sargent (2008) induce robust decision rules by replacing the mathematical expectations in (19) with risk-sensitivity operators. In particular, they substitute $(\mathbb{T}^1)(\theta_1)$ for E_{s^*, π^*} and replace E_z with $(\mathbb{T}^2)(\theta_2)$. This delivers a Bellman equation

$$(26) \quad V(s, \pi) = \max_v \left\{ r(s, v) + \mathbb{T}^2 \left[\mathbb{T}^1(\beta V(s^*, \pi^*)(s, v, \pi, z; \theta_1)) \right] (s, v, \pi; \theta_2) \right\}.$$

Notice that the parameters θ_1 and θ_2 are allowed 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. Cogley, Colacito, Hansen, and Sargent (2008) document how applications of these two operators have very different ramifications for experimentation in the context of their extended example that features competing conceptions of the Phillips curve.²⁵ Activating the \mathbb{T}^1 operator

²¹A change in the state dynamics would imply a misspecification in the evolution of the state probabilities.

²²In contrast Levin and Williams (2003) do not consider model averaging and implications for learning about which model fits the data better.

²³See Petersen, James, and Dupuis (2000) for an alternative approach to “structured uncertainty”.

²⁴See Taylor and Williams (2009) for a robustness comparison across alternative monetary policy rules.

²⁵When $\theta_1 = \theta_2$ the two operators applied in conjunction give the recursive formulation of risk sensitivity proposed in Hansen and Sargent (1995a), appropriately modified for the inclusion of hidden states.

reduces the value to experimentation because of the suspicions about the specifications of each model that are introduced. Activating the T^2 operator enhances the value to experimentation in order to reduce the ambiguity across models. Thus, the two notions of robustness embedded in these operators have offsetting impacts on the value of experimentation.

5.5. Sudden changes in beliefs. Hansen and Sargent (2008a) apply the T^1 and T^2 operators to build a model of sudden changes in expectations of long-run consumption growth ignited by news about consumption growth. Since the model envisions an endowment economy, the model is designed to focus on the impacts of beliefs on asset prices. Because concerns about robustness make a representative consumer especially averse to persistent uncertainty in consumption growth, fragile expectations created by model uncertainty transmit and induce what ordinary econometric procedures would measure as high and state-dependent market prices of risk.

Hansen and Sargent (2008a) analyze a setting in which there are two submodels of consumption growth. Let c_t be the logarithm of per capita consumption. Model $\iota \in \{0, 1\}$ has a more or less persistent component of consumption growth

$$\begin{aligned} c_{t+1} - c_t &= \mu(\iota) + z_t + \sigma_1(\iota)\varepsilon_{\iota,t+1} \\ z_{t+1}(\iota) &= \rho(\iota)z_t(\iota) + \sigma_2(\iota)\varepsilon_{2,t+1} \end{aligned}$$

where $\mu(\iota)$ is an unknown parameter with prior distribution $\mathcal{N}(\mu_c(\iota), \sigma_c(\iota))$, ε_t is an i.i.d. 2×1 vector process distributed $\mathcal{N}(0, I)$, and $z_0(\iota)$ is an unknown scalar distributed as $\mathcal{N}(\mu_x(\iota), \sigma_x(\iota))$. Model $\iota = 0$ has low $\rho(\iota)$ and makes consumption growth nearly i.i.d., while model $\iota = 1$ has $\rho(\iota)$ approaching 1, which, with a small value for $\sigma_2(\iota)$, gives consumption growth a highly persistent component of low conditional volatility but high unconditional volatility.

Bansal and Yaron (2004) tell us that these two models are difficult to distinguish using post World War II data for the United States. Hansen and Sargent (2008a) put an initial prior of .5 on these two submodels and calibrate the submodels so that the Bayesian posterior over the two submodels is .5 at the end of the sample. Thus, the two models are engineered so that the likelihood functions for the two submodels evaluated for the entire sample are identical. The solid blue line in figure 3 shows the Bayesian posterior on the long-run risk $\iota = 1$ model constructed in this way. Notice that while it wanders, it starts and ends at .5.

The higher green line shows the worst-case probability that emerges from applying a T^2 operator. The worst-case probabilities depicted in Figure 3 indicate that the representative consumer's concern for robustness makes him slant model selection probabilities towards the long-run risk model because, relative to the $\iota = 0$ model with less persistent consumption growth, the long-run risk $\iota = 1$ model has adverse consequences for discounted utility. A cautious investor mixes submodels by slanting probabilities towards the model with the lower discounted expected utility. Of especial interest in Figure 3 are recurrent episodes in which news expands the *gap*




FIGURE 3. Bayesian probability $\pi_t = E_t(\iota)$ attached to long-run risk model for growth in U.S. quarterly consumption (nondurables plus services) per capita for $p_0 = .5$ (lower line) and worst-case probability \check{p}_t (higher line). We have calibrated θ_1 to give a detection error probability conditional on observing $\mu(0), \mu(1)$ and z_t of .4 and θ_2 to give a detection error probability of .2 for the distribution of $c_{t+1} - c_t$.

between the worst-case probability and the Bayesian probability π_t assigned to the long-run risk model $\iota = 1$. This provides Hansen and Sargent (2008a) with a way to capture instability of beliefs alluded to by Keynes in the passage quoted above.

Hansen and Sargent (2008a) explain how the dynamics of continuation utilities conditioned on the two submodels contribute to countercyclical market prices of risk. The representative consumer regards an adverse shock to consumption growth as portending permanent bad news because he increases the worst-case probability \check{p}_t that he puts on the $\iota = 1$ long-run risk model, while he interprets a positive shock to consumption growth as only temporary good news because he raises the probability $1 - \check{p}_t$ that he attaches to the $\iota = 0$ model that has less persistent consumption growth. Thus, the representative consumer is pessimistic in interpreting good news as temporary and bad news as permanent.

5.6. Adaptive Models. In principle, the approach of the preceding sections could be applied to our basic linear-quadratic setting by positing a stochastic process of the A, B matrices so that there is a *tracking problem*. The decision maker must learn about a perpetually moving target. Current and past data must be used to make inferences about the process for the A, B matrices. But specifying the problem completely now becomes quite demanding, as the decision maker is compelled to take a stand on the stochastic evolution of the matrices A, B . The solutions are also much more difficult to compute because the decision maker at date t must deduce beliefs about the future trajectory of A, B given current and past information. The

greater demands on model specification may cause decision makers to second guess the reasonableness of the auxiliary assumptions that render the decision analysis tractable and credible. This leads us to discuss a non-Bayesian approach to tracking problems.

This approach to model uncertainty comes from distinct literatures on adaptive control and vector autoregressions with random coefficients.²⁶ What is sometimes called *passive* adaptive control is occasionally justified as providing robustness against parameter drift coming from model misspecification.

Thus, a random coefficients model captures doubts about the values of components of the matrices A, B by specifying that

$$x_{t+1} = A_t x_t + B_t u_t + C w_{t+1}$$

where $w_{t+1} \sim \mathcal{N}(0, I)$ and the coefficients are described by

$$(27) \quad \begin{bmatrix} \text{col}(A_{t+1}) \\ \text{col}(B_{t+1}) \end{bmatrix} = \begin{bmatrix} \text{col}(A_t) \\ \text{col}(B_t) \end{bmatrix} + \begin{bmatrix} \eta_{A,t+1} \\ \eta_{B,t+1} \end{bmatrix}$$

where now

$$\nu_{t+1} \equiv \begin{bmatrix} w_{t+1} \\ \eta_{A,t+1} \\ \eta_{B,t+1} \end{bmatrix}$$

is a vector of independently and identically distributed shocks with specified covariance matrix Q , and $\text{col}(A)$ is the vectorization of A . Assuming that the state x_t is observed at t , a decision maker could use a tracking algorithm

$$\begin{bmatrix} \text{col}(\hat{A}_{t+1}) \\ \text{col}(\hat{B}_{t+1}) \end{bmatrix} = \begin{bmatrix} \text{col}(\hat{A}_t) \\ \text{col}(\hat{B}_t) \end{bmatrix} + \gamma_t h(x_t, u_t, x_{t-1}; \text{col}(\hat{A}_t), \text{col}(\hat{B}_t)),$$

where γ_t is a ‘gain sequence’ and $h(\cdot)$ is a vector of time- t values of ‘sample orthogonality conditions’. For example, a least squares algorithm for estimating A, B would set $\gamma_t = \frac{1}{t}$. This would be a good algorithm if A, B were not time varying. When they are time-varying (i.e., some of the components of Q corresponding to A, B are not zero), it is better to set γ_t to a constant. This in effect discounts past observations.

Problem 5. (*Adaptive Control*)

To get what control theorists call an adaptive control model, or what Kreps (1998) calls an anticipated utility model, for each t solve the fixed point problem (4) subject to

$$(28) \quad x^* = \hat{A}_t x + \hat{B}_t u + C w^*.$$

²⁶See Kreps (1998) and Sargent (1999b) for related accounts of this approach. See Marcat and Nicolini (2003), Sargent, Williams, and Zha (2006), Sargent, Williams, and Zha (2009), and Carboni and Ellison (2009) for empirical applications.

The solution is a control law $u_t = -F_t x_t$ that depends on the most recent estimates of A, B through the solution of the Bellman equation (4).

The adaptive model misuses the Bellman equation (4), which is designed to be used under the assumption that the A, B matrices in the transition law are time-invariant. Our adaptive controller uses this marred procedure because he wants a workable procedure for updating his beliefs using past data and also for looking into the future while making decisions today. He is of two minds: when determining the control $u_t = -F_t x_t$ at t , he pretends that $(A, B) = (\hat{A}_t, \hat{B}_t)$ will remain fixed in the future; but each period when new data on the state x_t are revealed, he updates his estimates. This is not the procedure of a Bayesian who believes (27), as we have seen above. It is often excused because it is much simpler than a Bayesian analysis or some loosely defined kind of ‘bounded rationality’.

5.7. State prediction. Another way to incorporate learning in a tractable manner is to shift the focus from the transition law to the state. Suppose the decision maker is not able to observe the entire state vector and instead must make inferences about this vector. Since the state vector evolves over time, we have another variant of a tracking problem.

When a problem can be formulated as learning about an observed piece of the original state x_t , the construction of decision rules with and without concerns about robustness becomes tractable.²⁷ Suppose that the A, B, C matrices are known *a priori* but that some component of the state vector is not observed. Instead, the decision maker sees an observation vector y constructed from x

$$y = Sx.$$

While some combinations of x can be directly inferred from y , others cannot. Since the unobserved components of the state vector process x may be serially correlated, the history of y can help in making inferences about the current state.

Suppose, for instance, that in a consumption-savings problem, a consumer faces a stochastic process for labor income. This process might be directly observable, but it might have two components that cannot be disentangled: a permanent component and a transitory component. Past labor incomes will convey information about the magnitude of each of the components. This past information, however, will typically not reveal perfectly the permanent and transitory pieces. Figure 4 shows impulse response functions for the two components of the endowment process estimated by Hansen, Sargent, and Tallarini (1999). The first two panels display impulse responses for two orthogonal components of the endowment, one of which, d^1 , is estimated to resemble a permanent component, the other of which, d^2 is more transitory. The third panel shows the impulse response for the univariate (Wold) representation for the total endowment $d_t = d_t^1 + d_t^2$.

²⁷See Jovanovic (1979) and Jovanovic and Nyarko (1996) for examples of this idea.

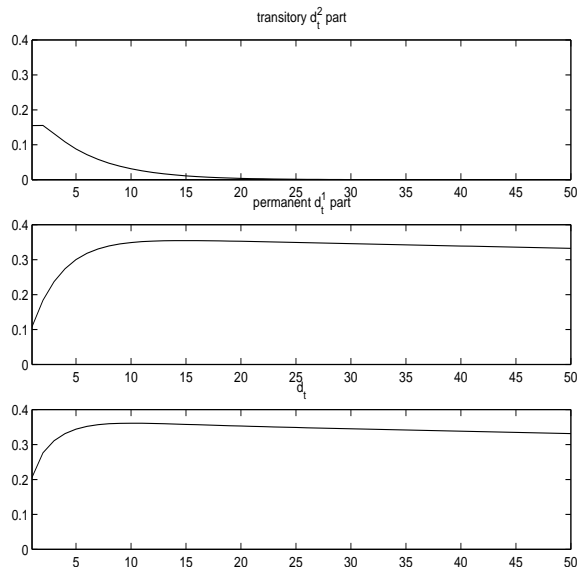


FIGURE 4. Impulse responses for two components of endowment process and their sum in a model of Hansen, Sargent, and Tallarini (1999). The top panel is the impulse response of the transitory component d^2 to an innovation in d^2 ; the middle panel, the impulse response of the permanent component d^1 to its innovation; the bottom panel is the impulse response of the sum $d_t = d_t^1 + d_t^2$ to its own innovation.

Figure 5 depicts the transitory and permanent components to income implied by the parameter estimates of Hansen, Sargent, and Tallarini (1999). Their model implies that the separate components d_t^i can be recovered *ex post* from the detrended data on consumption and investment that they used to estimate the parameters. Figure 6 uses Bayesian updating (Kalman filtering) to form estimators of d_t^1, d_t^2 assuming that the parameters of the two endowment processes are known, but that only the history of the *total* endowment d_t is observed at t . Note that these filtered estimates in Figure 6 are smoother than the actual components.

Alternatively, consider a stochastic growth model of the type advocated by Brock and Mirman (1972), but with a twist. Brock and Mirman studied the efficient evolution of capital in an environment in which there is a stochastic evolution for the technology shock. Consider a setup in which the technology shock has two components. Small shocks hit repeatedly over time and large technological shifts occur infrequently. The technology shifts alter the rate of technological progress. Investors may not be able to disentangle small repeated shifts from large but infrequent shifts




FIGURE 5. Actual permanent and transitory components of endowment process from Hansen, Sargent, Tallarini (1999) model.

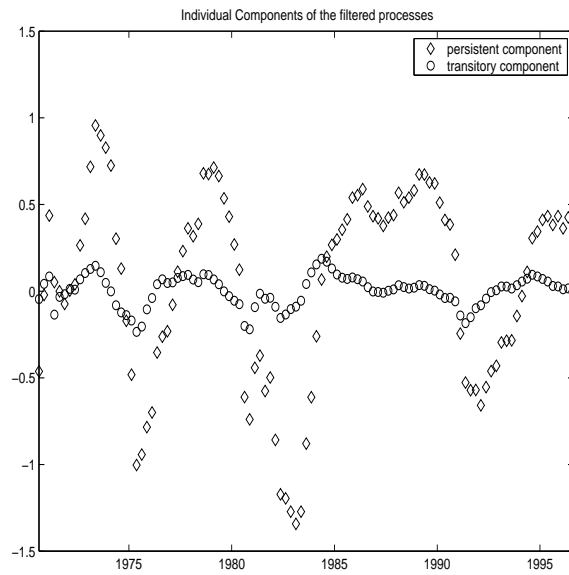


FIGURE 6. Filtered estimates of permanent and transitory components of endowment process from Hansen, Sargent, Tallarini (1999) model.

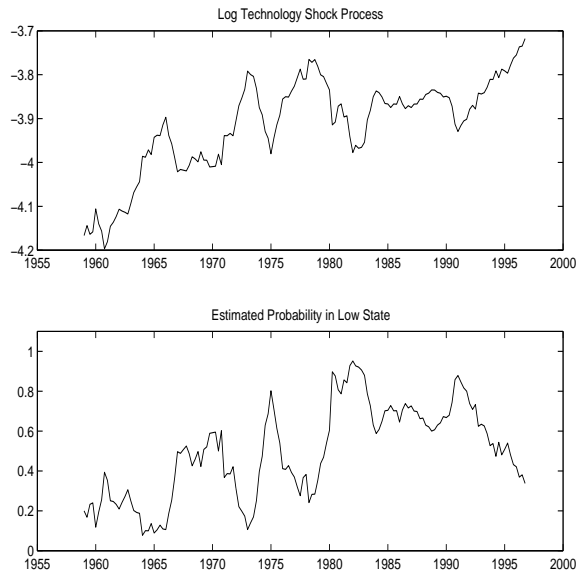


FIGURE 7. Top panel: the growth rate of the Solow residual, a measure of the rate of technological growth. Bottom panel: the probability that growth rate of the Solow residual is in the low-growth state.

in technological growth.²⁸ For example, investors may not have perfect information about the timing of a productivity slowdown that probably occurred in the seventies. Suppose investors look at the current and past levels of productivity to make inferences about whether technological growth is high or low. Repeated small shocks disguise the actual growth rate. Figure 7 reports the technology process extracted from postwar data and also shows the probabilities of being in a low growth state. Notice that during the so-called productivity slowdown of the seventies, even Bayesian learners would not be particularly confident in this classification for much of the time period. Learning about technological growth from historical data is potentially important in this setting.

5.8. The Kalman filter. Suppose for the moment that we abstract from concerns about robustness. In models with hidden state variables, there is a direct and elegant counterpart to the control solutions described above. It is called the Kalman filter, and recursively forms Bayesian forecasts of the current state vector given current and past information. Let \hat{x} denote the estimated state. In a stochastic counterpart

²⁸It is most convenient to model the growth rate shift as a jump process with a small number of states. See Cagetti, Hansen, Sargent, and Williams (2002) for an illustration. It is most convenient to formulate this problem in continuous time. The Markov jump component pushes us out of the realm of the linear models studied here.

to a steady state, the estimated state and the observed y^* evolve according to:

$$(29) \quad \hat{x}^* = A\hat{x} + Bu + G_x\hat{w}^*$$

$$(30) \quad y^* = SA\hat{x} + SBu + G_y\hat{w}^*$$

where G_y is nonsingular. While the matrices A and B are the same, the shocks are different, reflecting the smaller information set available to the decision maker. The nonsingularity of G_y guarantees that the new shock \hat{w} can be recovered from next-period's data y^* via the formula

$$(31) \quad \hat{w} = (G_y)^{-1}(y^* - SA\hat{x} - SBu).$$

However, the original w^* cannot generally be recovered from y^* . The Kalman filter delivers a new information state that is matched to the information set of a decision maker. In particular, it produces the matrices G_x and G_y .²⁹

In many decision problems confronted by macroeconomists, the target depends only on the observable component of the state, and thus:³⁰

$$(32) \quad z = H\hat{x} + Ju,$$

5.9. Ordinary filtering and control. With no preference for robustness, Bayesian learning has a modest impact on the decision problem (1).

Problem 6. (*Combined Control and Prediction*)

The steady-state Kalman filter produces a new state vector, state evolution equation (29) and target equation (32). These replace the original state evolution equation (1) and target equation (2). The G_x matrix replaces the C matrix, but because of certainty equivalence, this has no impact on the decision rule computation. The optimal control law is the same as in problem 1, but it is evaluated at the new (estimated) state \hat{x} generated recursively by the Kalman filter.

5.10. Robust filtering and control. To put a preference for robustness into the decision problem, we again introduce a second agent and formulate a dynamic recursive two-person game. We consider two such games. They differ in how the second agent can deceive the first agent.

In decision problems with only *terminal* rewards, it is known that Bayesian-Kalman filtering is robust for reasons that are subtle (see Basar and Bernhard (1995) chapter 7 and Hansen and Sargent (2008b), chapters 17 and 18, for discussions). Suppose the decision maker at date t has no concerns about past rewards. He only cares about rewards in current and future time periods. This decision maker will have data available from the past in making decisions. Bayesian updating using the Kalman filter remains a defensible way to use this past information, even if model

²⁹In fact, the matrices G_x and G_y are not unique but the so-called gain matrix $K = G_x(G_y)^{-1}$ is.

³⁰A more general problem in which z depends directly on hidden components of the state vector can also be handled.

misspecification is entertained. Control theorists break this result by having the decision maker continue to care about initial period targets even as time evolves (*e.g.* see Basar and Bernhard (1995) and Zhou, Doyle, and Glover (1996)). In the games posed below, we take a recursive perspective on preferences by having time t decision makers only care about current and future targets. That justifies our continued use of the Kalman filter even when there is model misspecification and delivers separation of prediction and control that is not present in the counterpart control theory literature. See Hansen and Sargent (2008b), Hansen, Sargent, and Wang (2002), and Cagetti, Hansen, Sargent, and Williams (2002) for an elaboration.

Game 7. (*Robust Control and Prediction, i*)

To compute a robust control law, we solve the two-person zero-sum game 3 but with the information or predicted state \hat{x} replacing the original state x . Since we perturb evolution equation (29) instead of (1), we substitute the matrix G_x for C when solving the robust control problem. Since the equilibrium of our earlier two-person zero-sum game depended on the matrix C , the matrix G_x produced by the Kalman filter alters the control law.

Except for replacing C by G_x and the unobserved state x with its predicted state \hat{x} , the equilibria of game 7 and game 3 coincide.³¹ The separation of estimation and control makes it easy to modify our previous analysis to accommodate unobserved states.

A complaint about game 7 is that the original state evolution was relegated to the background by forgetting the structure for which the innovations representation (29), (30) is an outcome. That is, when solving the robust control problem, we failed to consider direct perturbations in the evolution of the original state vector, and only explored indirect perturbations from the evolution of the predicted state. The premise underlying game 3 is that the state x is directly observable. When x is not observed, an information state \hat{x} is formed from past history, but x is not observed. Game 7 fails to take account of this distinction.

To formulate an alternative game that recognizes this distinction, we revert to the original state evolution equation:

$$x^* = Ax + Bu + Cw^*.$$

The state x is unknown, but can be predicted by current and past values of y using the Kalman filter. Substituting \hat{x} for x yields:

$$(33) \quad x^* = A\hat{x} + Bu + \check{G}\check{w}^*,$$

where \check{w}^* has an identity matrix as its covariance matrix and the (steady-state) forecast-error covariance matrix for x^* given current and past values of y is $\check{G}\check{G}'$.

³¹Although the matrix G_x is not unique, the implied covariance matrix $G_x(G_x)'$ is unique. The robust control depends on G_x only through the covariance matrix $G_x(G_x)'$.

To study robustness, we disguise the model misspecification by the shock \tilde{w}^* . Notice that the dimension of \tilde{w}^* is typically greater than the dimension of \hat{w}^* , providing more room for deception because we use the actual next-period state x^* on the left-hand side of evolution equation (33) instead of the constructed information state \hat{x}^* . Thus, we allow perturbations in the evolution of the unobserved state vector when entertaining model misspecification.

Game 8. (*Robust Control and Prediction, ii*)

To compute a robust control law, we solve the two-person zero-sum game 3 but with the matrix \check{G} used in place of C .

For a given choice of the robustness parameter θ , concern about misspecification will be more potent in game 8 than in the other two-person zero-sum games. Mechanically, this is because

$$\begin{aligned}\check{G}(\check{G})' &\geq CC' \\ \check{G}(\check{G})' &\geq G_x(G_x)'.\end{aligned}$$

The first inequality compares the covariance matrix of x^* conditioned on current and past values of y to the covariance matrix of x^* conditioned on the current state x . The second inequality compares the covariance of x^* to the covariance of its estimator \hat{x}^* , both conditioned on current and past values of y . These inequalities show that there is more latitude to hide model misspecification in game 8 than in the other two robustness games. The enlarged covariance structure makes statistical detection more challenging. The fact that the state is unobserved gives robustness more potency in game 8 than in game 3.³² The fact that the decision makers explore the evolution of x^* instead of the information state \hat{x}^* gives robustness more potency in game 8 than 7.³³

In summary, the elegant decision theory for combined control and prediction has direct extensions to accommodate robustness. Recursivity in decision-making makes Bayesian updating methods justifiable for making predictions while looking back at current and past data even when there are concerns about model misspecification. When making decisions that have future consequences, robust control techniques alter decision rules in much the same way as when the state vector is fully observed. These ideas are reflected in games 7 and 8.

³²Game 3 corresponds to the outcome in risk-sensitive joint filtering and control. See Whittle (1980). Thus, when filtering is part of the problem, the correspondence between risk-sensitive control and preferences for robustness is modified.

³³As emphasized in Hansen, Sargent, and Wang (2002), holding θ fixed across games is different than holding detection errors probabilities fixed. See Barillas, Hansen, and Sargent (2009) for an illustration of this in the context of an example that links risk-premia culled from asset prices to measuring the uncertainty costs associated with aggregate fluctuations.

5.11. Adaptive control versus robust control. The robustness of Bayesian updating is tied to the notion of an approximating model (A, B, C) and perturbations around that model. The adaptive control problem 5 is aimed at eliminating the commitment to a time-invariant benchmark model. While a more flexible view is adopted for prediction, a commitment to the *estimated* model is exploited in the design of a control law for reasons of tractability. Thus, robust control and prediction combines Bayesian learning (about an unknown state vector) with robust control, while adaptive control combines flexible learning about parameters with standard control methods.

6. ROBUSTNESS IN ACTION

6.1. Robustness in a simple macroeconomic model. We use Ball's model to illustrate the robustness attained by alternative settings of the parameter θ . For Ball's model, we present Figure 8 to show that while robust rules do less well when the approximating model actually generates the data, their performance deteriorates more slowly with departures of the data generating mechanism from the approximating model.

Following the risk-sensitive control literature, we transform θ into the risk-sensitivity parameter $\sigma \equiv -\theta^{-1}$. Figure 8 plots the value $-E(\pi^2 + y^2)$ attained by three rules under the the worst-case model for the value of σ on the ordinate axis. The rules are those for the three values $\sigma = 0, -.04, -.085$. Recall how the detection error probabilities computed above associate a value of $\theta = -0.085$ with a detection error probability of about .1. Notice how the robust rules (those computed with preference parameter $\sigma = -.04$ or $-.085$) have values that deteriorate at a lower rate with model misspecification (they are flatter). Notice that the rule for $\sigma = -.085$ does worse than the $\sigma = 0$ or $\sigma = -.04$ rules when $\sigma = 0$, but is more robust in deteriorating less when the model is misspecified. Next, we turn to various ways of characterizing the features that make the robust rules more robust.

6.2. Responsiveness. A common method for studying implications of dynamic economic models is to compute the *impulse responses* of economic variables to shocks. Formally, these responses are a sequence of dynamic multipliers that show how a shock vector w_t alters current and future values of the state vector x_t and the target z_t tomorrow. These same impulse response sequences provide insights into how concerns about robustness alter the decision-making process.

6.2.1. Impulse Responses. Let F be a candidate control law and suppose there is no model misspecification. Thus, the state vector x_t evolves according to:

$$x_{t+1} = (A - BF)x_t + Cw_{t+1}.$$

and the target is now given by:

$$z_t = (H - JF)x_t.$$




FIGURE 8. Value of $-E(\pi^2 + y^2)$ for three decision rules when the data are generated by the worst-case model associated with the value of σ on the horizontal axis: $\sigma = 0$ rule (solid line), $\sigma = -.04$ rule (dashed-dotted line), $\sigma = -.085$ (dashed) line.

To compute an impulse response sequence, we run the counterfactual experiment of setting x_{-1} to zero, w_0 to some arbitrary vector of numbers, and all future w_t 's to zero. It is straightforward to show that the resulting targets are:

$$(34) \quad z_t = (H - JF)(A - BF)^t C w_0.$$

The impulse response sequence is just the sequence of matrices: $\mathcal{I}(F, 0) = (H - JF)C$, $\mathcal{I}(F, 1) = (H - JF)(A - BF)C$, ... , $\mathcal{I}(F, t - 1) = (H - JF)(A - BF)^{t-1}C$,

Under this counterfactual experiment, the objective (3) is given by:

$$(35) \quad -\frac{1}{2}(w_0)' \sum_{t=0}^{\infty} \beta^t \mathcal{I}(F, t - 1)' \mathcal{I}(F, t - 1) w_0.$$

Of course, shocks occur in all periods not just period zero, so the actual object should take these into account as well. Since the shocks are presumed to be independent over time, the contributions of shocks at different time periods can effectively be uncoupled (see the discussion of spectral utility in Whiteman (1986)). Absorbing the discounting into the impulse responses, we see that in the absence of model misspecification, the aim of the decision maker is to choose F to make the sequence of matrices $\mathcal{I}(F, 0)$, $\sqrt{\beta}\mathcal{I}(F, 1)$, ..., $\sqrt{\beta^t}\mathcal{I}(F, t)$, ... small in magnitude. Thus, (35) induces no preferences over specific patterns of the impulse response sequence, only

about the overall *magnitude* of the sequence as measured by the discounted *sum* (35).

Even though we have only considered a degenerate shock sequence, maximizing objective (3) by choice of F gives precisely the solution to problem 1. In particular, the optimal control law does not depend on the choice of w_0 for $w_0 \neq 0$. We summarize this in:

Claim 9. (*Frequency Domain Problem*) *For every w_0 , the solution of the problem of choosing a fixed F to maximize (35) is the same \hat{F} that solves problem (1). This problem induces no preferences about the shape of the impulse response function, only about its magnitude as measured by (35).*

In the next subsection, we shall see that a preference for robustness induces preferences about the shape of the impulse response function as well as its magnitude.

6.2.2. *Model misspecification with filtering.* Consider now potential model misspecification. As in game 3, we introduce a second, minimizing agent. In our counterfactual experiment, suppose this second agent can choose future v_t 's to damage the performance of the decision rule F . Thus, under our hypothetical experiment, we envision state and target equations:

$$\begin{aligned}x_{t+1} &= Ax_t + Bu_t + Cv_t \\z_t &= Hx_t + Jv_t\end{aligned}$$

with $x_0 = Cw_0$. By conditioning on an initial w_0 , we are free to think of the second agent as choosing a sequence of the v_t 's that might depend on the initial w_0 . A given v_t will influence current and future targets via the impulse response sequence derived above.

To limit the damage caused by the malevolent agent, we penalize the choice of the v_t sequence by using the robustness multiplier parameter θ . Thus, the non-recursive objective for the two-player zero-sum dynamic game is:

$$(36) \quad - \sum_{t=0}^{\infty} \beta^t \{ |z_t| - \theta |v_t|^2 \}.$$

When the robustness parameter θ is large, the implicit constraint on the magnitude of the sequence of v_t 's is small and very little model misspecification is tolerated. Smaller values of θ permit sequences v_t that are larger in magnitude. A malevolent player agent chooses a v_t sequence to minimize (36) To construct a robust control law, the original decision maker then maximizes (36) by choice of F . This non-recursive representation of the game can be solved using the Fourier transform techniques employed by Whiteman (1986), Kasa (1999), and Christiano and Fitzgerald (1998). See Hansen and Sargent (2008b), chapter 8, for a formal development. This non-recursive game has the same solution as the recursive game 3.

Before describing some details, it is easy to describe informally how the malevolent agent will behave. He will detect seasonal, cyclical or long-run patterns in the implied impulse response sequences $\{\sqrt{\beta}I(F, t)\}_{t=0}^{\infty}$, then use his limited resources to concentrate deception at those frequencies. Thus, the minimizing agent will make the v_t 's have cyclical components at those frequencies in the impulse response function at which the maximizing agent's choice of F leaves himself most vulnerable as measured by (35).

Here the mathematical tool of Fourier transforms allows us to summarize the impulse response function in the frequency domain.³⁴ Imagine using a representation of the components of the specification error v_t sequence in terms of sines and cosines to investigate the effects on the objective function when misspecification is confined to particular frequencies. Searching over frequencies for the most damaging effects on the objective allows the minimizing agent to put particular temporal patterns into the v_t 's. It is necessary to view the composite contribution of entire v_t sequence, including its temporal pattern.

An impulse response sequence summarizes how future targets respond to a current period v_t ; a Fourier transform of the impulse response function quantifies how future targets respond to v_t sequences that are pure cosine waves. When the minimizing agent chooses a temporally dependent v_t sequence, the maximizing agent should care about the temporal pattern of the impulse response sequence, not just its overall magnitude.³⁵ The minimizing agent in general will find that some particular frequencies (e.g., a cosine wave of given frequency for the v_t 's) will most efficiently exploit model misspecification. To respond best to the actions of the minimizing agent, therefore, in addition to making the impulse response sequence small, now the maximizing agent wants to design a control law F in part to flatten the *frequency sensitivity* of the (appropriately discounted) impulse response sequence. This concern causes a trade-off across frequencies to emerge. The robustness parameter θ balances a tension between asking that impulse responses be small in magnitude and also that they be insensitive to model misspecification.

6.3. Some frequency domain details. To investigate these ideas in more detail, we use some arithmetic of complex numbers. Recall that

$$\exp(i\omega t) = \cos(\omega t) + i \sin(\omega t).$$

We can extract a frequency component from the misspecification sequence $\{v_t\}$ using a Fourier transform. Define:

$$\mathcal{FT}(v)(\omega) = \sum_{t=0}^{\infty} \beta^{t/2} v_t \exp(i\omega t), \quad \omega \in [-\pi, \pi].$$

³⁴Also see Brock, Durlauf, and Rondina (2008).

³⁵It was the absence of the temporal dependence in the v_t 's under the approximating model that left the maximizing agent indifferent to the shape of the impulse response function in (35).

We can interpret

$$\mathcal{FT}(v)(\omega) \exp(-i\omega t)$$

as the frequency ω component of the misspecification sequence. Our justification for this claim comes from the integration recovery (or inversion) formula:

$$\beta^{t/2} v_t = \frac{1}{2\pi} \int_{-\pi}^{\pi} \mathcal{FT}(v)(\omega) \exp(-i\omega t) d\omega.$$

Thus, we have an additive decomposition over the frequency components. By adding up or integrating over these frequencies, we recover the misspecification sequence in the time domain. Moreover, the squared magnitude of the misspecification sequence can be depicted as an integral:

$$\sum_{t=0}^{\infty} \beta^t v_t \cdot v_t = \frac{1}{2\pi} \int_{-\pi}^{\pi} |\mathcal{FT}(v)(\omega)|^2 d\omega$$

Thus, Fourier transforms provide a convenient toolkit for thinking formally about misspecification in terms of frequency decompositions.

It may appear troubling that the frequency components are complex. However, by combining contribution at frequencies ω and $-\omega$, we obtain sequences of real vectors. The periodicity of frequency ω and frequency $-\omega$ are identical, so it makes sense to treat these two components as a composite contribution. Moreover, $|\mathcal{FT}(v)(\omega)| = |\mathcal{FT}(v)(-\omega)|$.

We can get a version of this decomposition for the appropriately discounted target vector sequence.³⁶ This calculation results in the following formula for the Fourier transform $\mathcal{FT}(z)(\omega)$ of the ‘target’ z_t sequence:

$$\mathcal{FT}(z)(\omega) = h(\omega)[w_0 + \exp(i\omega)\mathcal{FT}(v)(\omega)]$$

where the matrix function

$$\begin{aligned} h(\omega) &= (H - JF)[I - \sqrt{\beta}(A - BF) \exp(i\omega)]^{-1}C \\ &= \sum_{t=0}^{\infty} \beta^{t/2} \mathcal{I}(F, t) \exp(i\omega t). \end{aligned}$$

is the Fourier transform of the sequence of impulse responses from the shocks to the target z_t . This transform depends implicitly on the choice of control law F . This Fourier transform describes how frequency components of the misspecification sequence influence the corresponding frequency components of the target sequence. When the matrix $h(\omega)$ is large in magnitude relative to other frequencies, frequency ω is particularly vulnerable to misspecification.

³⁶That cosine shocks lead to cosine responses of the same frequency reflects the linearity of the model. In nonlinear models, the response to a cosine wave shock is more complicated.

Objective (36) has a frequency representation given by:

$$-\frac{1}{4\pi} \int_{-\pi}^{\pi} (|\mathcal{FT}(z)(\omega)|^2 - \theta |\mathcal{FT}(v)(\omega)|^2) d\omega.$$

The malevolent agent chooses to minimize this objective by choice of $\mathcal{FT}(v)(\omega)$. The control law F is then chosen to minimize the objective. As established in Hansen and Sargent (2008b), chapter 8, this is equivalent to ranking control laws F using the frequency-based *entropy* criterion:

$$(37) \quad \text{entropy} = -\frac{1}{2\pi} \int_{-\pi}^{\pi} \log \det[\theta I - h(\omega)'h(-\omega)] d\omega.$$

See Hansen and Sargent (2008b) for an explanation of how this criterion induces the same preferences over decision rules F as the two-player game 3. Lowering θ causes the decision maker to design F_θ to make (trace $h(\omega)'h(-\omega)$) flatter as a function of frequency, thereby lowering its larger values at the cost of raising smaller ones. Flattening (trace $h(\omega)'h(-\omega)$) makes the realized value of the criterion function less sensitive to departures of the shocks from the benchmark specification of no serial correlation.

6.3.1. *A limiting version of robustness.* There are limits on the size of the robustness parameter θ . When θ is too small, it is known that the two-player zero-sum game suffers a *breakdown*. The fictitious malevolent player can inflict sufficient damage that the objective function remains at $-\infty$ regardless of the control law F . The critical value of θ can be found by solving:

$$\underline{\theta} = \sup_v \frac{1}{2\pi} \int_{-\pi}^{\pi} |h(\omega)\mathcal{FT}(v)(\omega)|^2 d\omega$$

subject to

$$\frac{1}{2\pi} \int_{-\pi}^{\pi} |\mathcal{FT}(v)(\omega)|^2 d\omega = 1.$$

The sup is typically not attained, but is approximated by a sequence that isolates one particular frequency.

The critical value $\underline{\theta}$ depends on the choice of control law F . One (somewhat extreme) version of robust control theory, called H_∞ control theory, instructs a decision maker to select a control law to make this critical value of θ as small as possible.

6.3.2. *A related econometric defense for filtering.* In econometric analyses, it is often argued that time series data should be filtered before estimation to avoid contaminating parameters. Indeed, frequency decompositions can be used to justify such methods. The method called *spectral analysis* is about decomposing time series into frequency components. Consider an econometrician with a formal economic model to be estimated. He suspects, however, that the model may not be well suited to

explain all of the component movements in the time series. For instance, many macroeconomic models are not well designed to explain seasonal frequencies. The same is sometimes claimed for low frequency movements as well. In this sense the data may be *contaminated vis a vis* the underlying economic model.³⁷

One solution to this problem would be to put a prior distribution over all possible forms of contamination and to form a hyper model by integrating over this contamination. As we have argued above, that removes concerns about model misspecification from discussion, but arguably in a contrived way. Also, this approach will not give rise to the common applied method of filtering the data to eliminate particular frequencies where the most misspecification is suspected.

Alternatively, we could formalize the suspicion of data contamination by introducing a malevolent agent who has the ability to contaminate time series data over some frequency range, say seasonal frequencies or low frequencies that correspond to long-run movements in the time series. This contamination can undermine parameter estimation in a way formalized in the frequency domain by Sims (1972) for least squares regression models and Sims (1993) and Hansen and Sargent (1993) for multivariate time-series models. Sims (1974) and Wallis (1974) used frequency domain characterizations to justify a seasonal-adjustment filter and to provide guidance about the appropriate structure of the filter. They found that if one suspects that a model is better specified at some frequencies than others, then it makes sense to diminish approximation errors by filtering the data to eliminate frequencies most vulnerable to misspecification.

Consider a two-player zero-sum game to formulate this defense. If an econometrician suspects that a model is better specified at some frequencies than others, this can be operationalized by allowing the malevolent agent to concentrate his mischief-making only at those frequencies, like the malevolent agent from robust control theory. The data filter used by the econometrician can emerge as a solution to an analogous two-player game. To arrest the effects of such mischief-making, the econometrician will design a filter to eliminate those frequencies from estimation.

Such an analysis provides a way to think about both seasonal adjustment and trend removal. Both can be regarded as procedures that remove frequency components with high power with the aim of focusing empirical analysis on frequencies where a model is better specified. Sims (1993) and Hansen and Sargent (1993) describe situations in which the cross-equation restrictions of misspecified rational expectations models provide better estimates of preference and technological parameters with seasonally adjusted data.

6.3.3. *Comparisons.* It is useful to compare the frequency domain analysis of data filtering with the frequency domain analysis of robust decision making. The robust decision maker achieves a robust rule by damping the influence of frequencies most vulnerable to misspecification. In the Sims (1993) analysis of data filtering, an

³⁷Or should we say that the model is *contaminated vis a vis* the data?

econometrician who fears misspecification and knows the approximation criterion is advised to choose a data filtering scheme that down plays frequencies at which he suspects the most misspecification. He does ‘window carpentry’ in crafting a filter to minimize the impact of specification error on the parameters estimates that he cares about.

6.4. Friedman: long and variable lags. We now return to Friedman’s concern about the use of misspecified models in the design of macroeconomic policies, and in particular to his view that lags in the effects of monetary policy are long and variable. The game theoretic formulation of robustness gives one possible expression to this concern about long and variable lags. That the lags are *long* is determined by the specification of the approximating model. (We will soon give an example in the form of the model of Laurence Ball.) That the lags are *variable* is captured by the innovation mean distortions v_t that are permitted to feed back arbitrarily on the history of states and controls. By representing misspecified dynamics, the v_t ’s can capture one sense of variable lags. Indeed, in the game theoretic construction of a robust rule, the decision maker acts as though he believes that the way that the worst-case v_{t+1} process feeds back on the state depends on his choice of decision rule F . This dependence can be expressed in the frequency domain in the way we have described. The structure of the original model (A, B, C) and the hypothetical control law F dictate which frequencies are most vulnerable to model misspecification. They might be low frequencies, as in Friedman’s celebrated permanent income model, or they might be business cycle or seasonal frequencies. Robust control laws are designed in part to dampen the impact of frequency responses induced by the v_t ’s. To blunt the role of this second player, under robustness the original player will aim to diminish the importance of the impulse response sequence beyond the initial response. The resulting control laws often lead to impulse responses that are greater at impact and are more muted in the tails. We give an illustration in the next subsection.

6.4.1. Robustness in Ball’s model. We return to Ball’s model and use it to illustrate how concerns about robustness affect frequency domain representations of impulse response functions. We discount the return function in Ball’s model, altering the object that the government would like to maximize to be

$$-E \sum_{t=0}^{\infty} \beta^t (\pi_t^2 + y_t^2).$$

We derive the associated robust rules for three values of the robustness parameter θ . In the frequency domain, the criterion can be represented as

$$H_2 = - \int_{-\pi}^{\pi} \text{trace}[h(\omega)'h(-\omega)]d\omega.$$




FIGURE 9. Frequency decompositions of trace $[h(\omega)'h(-\omega)]$ for objective function of Ball's model under three decision rules; discount factor $\beta = 1$ on left panel, $\beta = .9$ on right panel.

Here $h(\omega)$ is the transfer function from the shocks in Ball's model to the targets, the inflation rate and output. The transfer function h depends on the government's choice of a feedback rule F_θ . Ball computed F_∞ .

Figure 9 displays frequency decompositions of $[\text{trace } h(\omega)'h(-\omega)]$ for robust rules with $\beta = 1$ and $\beta = .9$. Figure 9 shows frequency domain decompositions of a government's objective function for three alternative policy rules labeled $\theta = +\infty, \theta = 10, \theta = 5$. The parameter θ measures a concern about robustness, with $\theta = +\infty$ corresponding to no concern about robustness, and lower values of θ representing a concern for misspecification. Of the three rules whose transfer functions are depicted in Figure 9, Ball's rule ($\theta = +\infty$) is the best under the approximating model because the area under the curve is the smallest.

The transfer function h gives a frequency domain representation of how targets respond to serially uncorrelated shocks. The frequency domain decomposition C depicted by the $\theta = +\infty$ curve in Figure 9 exposes the frequencies that are most vulnerable to small misspecifications of the temporal and feedback properties of the shocks. Low frequency misspecifications are most troublesome under Ball's optimal feedback rule because for those frequencies, $\text{trace}[h(\zeta)'h(\zeta)]$ is highest.

We can obtain more robust rules by optimizing the entropy criterion (37). Flattening the frequency response $\text{trace}[h(\omega)'h(-\omega)]$ is achieved by making the interest rate more sensitive to both y and e ; as we reduce θ , both a and b increase in the




FIGURE 10. Top panel: impulse responses of inflation to the shock η_t for three values of θ : $\theta = +\infty$ (solid line), $\theta = 10$ (dashed-dotted line), and $\theta = 5$ (dotted line), with $\beta = 1$. Bottom panel: impulse response of inflation to shock ϵ_t under same three values of θ .

feedback rule $r_t = ay_t + b\pi_t$.³⁸ This effect of activating a preference for robust rules has the following interpretation. Ball's model specifies that the shocks in (12),(13), (14) are serially uncorrelated. The no-concern about robustness $\theta = +\infty$ rule exposes the policy maker to the biggest costs if the shocks instead are actually highly positively serially correlated. This means that a policy maker who is worried about misspecification is most concerned about misreading what is actually a 'permanent' or 'long-lived' shock as a temporary (i.e., serially uncorrelated) one. To protect himself, the policy maker responds to serially uncorrelated shocks (under the approximating model) as though they were positively serially correlated. This response manifests itself in his making the interest rate more responsive to both y_t and π_t .

An interesting aspect of the two panels of Figure 9 is that in terms of trace $[h(\omega)'h(-\omega)]$, lowering the discount factor β has similar effects as lowering θ (compare the $\theta = 5$ curves in the two panels). Hansen, Sargent, and Tallarini (1999) uncovered a similar pattern in a permanent income model; they showed that there existed offsetting changes in β and θ that would leave the quantities (but not the prices) of a permanent income model unchanged.

Figure 10 displays impulse response functions of inflation to η_t (the shock in the Phillips curve) and ϵ_t (the shock in the IS curve) under the robust rules for $\theta = +\infty, 10, 5$ when $\beta = 1$. The panels show that activating preferences for robustness

³⁸See Sargent (1999a) for a discussion.

causes the impulse responses to damp out more quickly, which is consistent with the flatter trace $[h(\omega)'h(-\omega)]$ functions observed as we accentuate the preference for robustness. Note also that the impact effect of ϵ_t on inflation is increased with an increased preference for robustness.

6.5. Precaution. A property or limitation of the linear-quadratic decision problem 1 in the absence of robustness is that it displays certainty equivalence. The optimal decision rule does not depend on the matrix C that governs how shocks impinge on the the state evolution. The decision rule fails to adjust to presence of fluctuations induced by shocks (even though the decisions themselves do depend on the shocks). The rule would be the same even if shocks were set to zero. Thus, there is no motive for precaution.

The celebrated permanent-income model of Friedman (1956) (see Zeldes (1989) for an elaboration) has been criticized because it precludes a precautionary motive for savings. Leland (1968) and Miller (1974) extended Friedman's analysis to accommodate precautionary savings by moving outside the linear-quadratic functional forms given in problem 1. Notice that in decision problem 1, both the time t contribution to the objective function and the value function are quadratic and hence have zero third derivatives. For general decision problems under correct model specification, Kimball (1990) constructs a measure of precaution in terms of the third derivatives of the utility function or value function.

We have seen how a preference for robustness prompts the C matrix to influence behavior even within the confines of decision problem 1, which because it has a quadratic value function precludes a precautionary motive under correct model misspecification. Thus, a concern about model misspecification introduces an additional motivation for precaution beyond that suggested by Leland (1968) and Miller (1974). Shock variances play a role in this new mechanism because the model misspecification must be disguised to a statistician. Hansen, Sargent, and Tallarini (1999) are able to reinterpret Friedman's permanent income model of consumption as one in which the consumer is concerned about model misspecification. Under the robust interpretation, consumers discount the future more than under the certainty-equivalent interpretation. In spite of this discounting, consumers save in part because of concerns that their model of the stochastic evolution of income might be incorrect.

This new mechanism for precaution remains when robustness is introduced into the models studied by Leland (1968), Miller (1974), Kimball (1990), and others. In contrast to the precautionary behavior under correct model specification, robustness makes precaution depend on more than just third derivatives of value functions. The robust counterpart to Kimball (1990)'s measures of precaution depends on the *lower* order derivatives as well. This dependence on lower-order derivatives of the value function makes robust notions of precaution distinct from and potentially more

potent than the earlier notion of precaution coming from a non-zero third derivative of a value function.

6.6. Risk aversion. Economists are often perplexed by behavior of market participants that seems to indicate extreme risk aversion, for example, the behavior of asset prices and returns. To study risk aversion, economists want to confront decision makers with gambles described by known probabilities. From knowledge or guesses about how people would behave when confronted with specific and well defined risks, economists infer degrees of risk aversion that are *reasonable*. For instance, Barsky, Juster, Kimball, and Shapiro (1997) administered survey questions eliciting from people their willingness to participate in gambles. A distinct source of information about risk-aversion comes from measurements of risk-return trade-offs from financial market data. The implied connection between risk aversion as modeled by a preference parameter and risk-return trade-offs measured by financial econometricians was delineated by Hansen and Jagannathan (1991) and Cochrane and Hansen (1992). But evidence extracted in this way from historical security market data suggest that risk aversion implied by security market data is very much larger than that elicited from participants facing those hypothetical gambles with well understood probabilities.

There is a variety of responses to this discrepancy. One questions the appropriateness of extrapolating measures of risk-aversion extracted from hypothetical small gambles to much larger ones. For example, it has been claimed that people look more risk averse when facing small gambles than large ones (see Segal and Spivak (1990), Epstein and Melino (1995), and Rabin (1999)). Others question the empirical measurements of the risk-return trade-off because, for example, mean returns on equity are known to be difficult to measure reliably. Our statistical notion of robustness easily makes contact with such responses. Thus, a concern about robustness comes into play when agents believe that their probabilistic descriptions of risk might be misspecified. In security markets, precise quantification of risks is difficult. It turns out that there is a formal sense in which a preference for robustness as modeled above can be reinterpreted in terms of a large degree of risk aversion, treating the approximating model as known. This formal equivalence has manifestations in both decision-making and in prices. The observationally equivalent risk-averse or *risk-sensitive* interpretation of robust decision making was first provided by Jacobson (1973), but outside the recursive framework used here. Hansen and Sargent (1995b) build on work of Jacobson (1973) and Whittle (1980) to establish an equivalence between a preference for robustness and risk-sensitive preferences for the two-person zero-sum game 3. Anderson, Hansen, and Sargent (2003) and Hansen, Sargent, Turmuhambetova, and Williams (2006) extend this equivalence result to a larger class of recursive two-person zero-sum games. Thus, the decision rules that emerge from

robustness games are identical with those rules that come from *risk-sensitive* control problems with correctly specified models.³⁹

Hansen, Sargent, and Tallarini (1999), Tallarini (2000), Cagetti, Hansen, Sargent, and Williams (2002) show that in a class of stochastic growth models the effects of a preference for robustness or of a risk-sensitive adjustment to preferences are very difficult or impossible to detect in the behavior of quantities along, for example, aggregate data on consumption and investment. The reason is that in these models altering a preference for robustness has effects on quantities much like those that occur under a change in a discount factor. Alterations in the parameter measuring preference for robustness can be offset by a change in the discount factor and thereby leave consumption and investment allocations virtually unchanged.

However, that kind of observational equivalence result does not extend to asset prices. The same adjustments to preferences for robustness and discount factors that leave consumption and investment allocations unaltered can have marked effects on the value function of a planner in a representative agent economy and therefore on equilibrium market prices of risk. Hansen, Sargent, and Tallarini (1999) and Hansen, Sargent, and Wang (2002) have used this observation to study the effects of a preference for robustness on the theoretical value of the equity premium.

A simple and pedagogically convenient model of asset prices is obtained by studying the shadow prices from optimal resource allocation problems. These shadow prices contain a convenient decomposition of the risk-return trade-off. Let γ_t denote a vector of factor loadings, so that under an approximating model, the unpredictable component of the return is $\gamma_t \cdot w_{t+1}$. Let r_t^f denote the risk-free interest rate. Then the required mean return μ_t satisfies the factor pricing relation

$$\mu_t - r_t^f = \gamma_t \cdot q_t$$

where q_t is a vector of what are commonly referred to as factor risk prices. Changing the *price vector* q_t changes the required mean return. Economic models with risk-averse investors imply a specific *shadow price* formula for q_t . This formula depends explicitly on the risk preferences of the consumer. An implication of many economic models is that the magnitude $|q_t|$ of the price vector implied by a reasonable amount of risk aversion is too small to match empirical observations.

Introducing robustness gives us an additive decomposition for q_t in corresponding continuous-time models, as demonstrated by Anderson, Hansen, and Sargent (1999), Anderson, Hansen, and Sargent (2003) and Chen and Epstein (1998). One

³⁹This observational equivalence applies within an economy for perturbations modeled in the manner described here. It can be broken by either restricting the class of perturbations, by introducing differential penalty terms, or in some of formulations with hidden states. Also, this equivalence result applies for a given economic environment. The robustness penalty parameter θ should not be thought of as being invariant across environments with different state equations. Recall that in our discussion of calibration, we used specific aspects of the environment to constrain the magnitude of the penalty parameter.




FIGURE 11. Four-Period Market Price of Knightian Uncertainty versus Detection Error Probability for three models: HST denotes the model of Hansen, Sargent, and Tallarini (1999); ‘benchmark’ denotes their model modified along the lines of the first robust filtering game 7; ‘HSW’ denotes their model modified according to the second robust filtering game 8.

component is an explicit risk component and the other is a model uncertainty component. The model uncertainty component relates directly to the detection-error rates that emerge from the statistical discrimination problem described above. By exploiting this connection, Anderson, Hansen, and Sargent (2003) argue that it is reasonable to assign about a third of the observed $|q_t|$ to concerns about robustness. This interpretation is premised on the notion that the market *experiment* is fundamentally more complicated than the stylized experiments confronting people with well understood risks that are typically used to calibrate risk aversion. Faced with this complication, investors use models as approximations and make conservative adjustments. These adjustments show up prominently in security market prices even when they are disguised in macroeconomic aggregates.

Figure 11 is from Hansen, Sargent, and Wang (2002), who studied the contribution to the market price of risk from a concern about robustness in three models: the basic model of Hansen, Sargent, and Tallarini (1999) and two modified versions of it in which agents do not observe the state and so must filter. Those two versions corresponded to the two robust filtering games 7 and 8 described above. Figure 11 graphs the contribution to the market price of risk of four-period securities coming

from robustness for each of these models graphed against the detection error probability. Freezing the detection error probability across models makes the value of θ depend on the model. (See the discussion above about how the detection error probability depends on θ and the particular model.) Figure 11 affirms the tight link between detection error probabilities and the contribution of a concern about robustness to the market price of risk that was asserted by Anderson, Hansen, and Sargent (2003). Notice how the relationship between detection error probabilities and the contribution of robustness to the market price of risk does not depend on which model is selected. The figure also conveys that a preference for robustness corresponding to a plausible value of the detection error probability gives a substantial boost to the market price of risk.

7. CONCLUDING REMARKS

This paper has discussed work designed to account for a preference for decisions that are robust to model misspecification. We have focused mainly on single-agent decision problems. The decision maker evaluates decision rules against a set of models near his approximating model, and uses a two-person zero-sum game in which a malevolent agent chooses the model as an instrument to achieve robustness across the set of models.

We have not touched issues that arise in contexts where multiple agents want robustness. Those issues deserve serious attention. One issue is the appropriate equilibrium concept with multiple agents who fear model misspecification. We need an equilibrium concept to replace rational expectations. Hansen and Sargent (2008b), chapters 15 and 16, and Karantounias (2009) use an equilibrium concept that seems a natural extension of rational expectations because all agents share the same approximating model. Suitably viewed, the communism of models seen in rational expectations models extends only partially to this setting: now agents share an approximating model, but not necessarily their sets of surrounding models against which they value robustness, nor the synthesized worst-case models that they use to attain robustness. Anderson (2005) studies a pure endowment economy whose agents have what we would interpret as different concerns about robustness, and shows how the distribution of wealth over time is affected by those concerns.⁴⁰ Hansen and Sargent (2008b) (chapter 16), Kasa (1999), and Karantounias (2009) describe multi-agent problems in the form of Ramsey problems for a government facing a competitive private sector.

Preferences for robustness also bear on the Lucas critique. Lucas's critique is the assertion that rational expectations models make decision rules functions of stochastic processes of shocks and other variables exogenous to decision makers. To

⁴⁰Anderson embraces the risk-sensitivity interpretation of his preference specification, but it is also susceptible to a robustness interpretation. He studies a Pareto problem of a planner who shares the approximating model and recognizes the differing preferences of the agents.

each shock process, a rational expectations theory associates a distinct decision rule. Lucas criticized earlier work for violating this principle. What about robust decision theory? It partially affirms but partially belies the Lucas critique. For a given preference for robustness (that is, for a given $\theta < +\infty$), a distinct decision rule is associated with each approximating model, respecting the Lucas critique. However, for a given preference for robustness and a fixed approximating model, the decision maker is supposed to use the same decision rule for a *set* of models surrounding the approximating model, superficially ‘violating the Lucas critique’. Presumably, the decision maker would defend that violation by appealing to detection error probabilities large enough to make members of that set of models difficult to distinguish from the approximating model based on the data available.

APPENDIX A. GENERALIZATIONS

This appendix describes how the linear-quadratic setups in much of the text link to more general non-linear, non-Gaussian problems. We define relative entropy and how it relates to the term $v'_t v_t$ that plays such a vital role in the robust control problems treated in the text.

A.1. Relative entropy and multiplier problem. Let $V(\epsilon)$ be a (value) function of a random vector ϵ with density $\phi(\epsilon)$. Let $\theta > 0$ be a scalar penalty parameter. Consider a distorted density $\hat{\phi}(\epsilon) = m(\epsilon)\phi(\epsilon)$ where $m(\epsilon) \geq 0$ is evidently a likelihood ratio. The risk-sensitivity operator is defined in terms of the indirect utility function $\mathbb{T}V$ that emerges from:

Problem 10.

$$(38) \quad \mathbb{T}V = \min_{m(\epsilon) \geq 0} \int m(\epsilon) [V(\epsilon) + \theta \log m(\epsilon)] \phi(\epsilon) d\epsilon$$

subject to

$$(39) \quad \int m(\epsilon) \phi(\epsilon) d\epsilon = 1$$

Here $\int m(\epsilon) \log m(\epsilon) \phi(\epsilon) d\epsilon = \int \log m(\epsilon) \hat{\phi}(\epsilon) d\epsilon$ is the entropy of $\hat{\phi}$ relative to ϕ . The minimizing value of $m(\epsilon)$ is

$$(40) \quad m(\epsilon) = \frac{\exp(-V(\epsilon)/\theta)}{\int \exp(-V(\tilde{\epsilon})/\theta) \phi(\tilde{\epsilon}) d\tilde{\epsilon}}$$

and the indirect utility function satisfies

$$(41) \quad \mathbb{T}V = -\theta \log \int \exp(-V(\epsilon)/\theta) \phi(\epsilon) d\epsilon.$$

A.2. Relative entropy and Gaussian distributions. It is useful first to compute relative entropy for the case that ϕ is $\mathcal{N}(0, I)$ and $\hat{\phi}$ is $\mathcal{N}(w, \Sigma)$, where the covariance matrix Σ is nonsingular. We seek a formula for $\int m(\varepsilon) \log m(\varepsilon) \phi(\varepsilon) d\varepsilon = \int (\log \hat{\phi}(\varepsilon) - \log \phi(\varepsilon)) \hat{\phi}(\varepsilon) d\varepsilon$. The log-likelihood ratio is

$$(42) \quad \log \hat{\phi}(\varepsilon) - \log \phi(\varepsilon) = \frac{1}{2} [-(\varepsilon - w)' \Sigma^{-1} (\varepsilon - w) + \varepsilon' \varepsilon - \log \det \Sigma] .$$

Observe that

$$- \int \frac{1}{2} (\varepsilon - w)' \Sigma^{-1} (\varepsilon - w) \hat{\phi}(\varepsilon) d\varepsilon = -\frac{1}{2} \text{trace}(I).$$

Applying the identity $\varepsilon = w + (\varepsilon - w)$ gives

$$\frac{1}{2} \varepsilon' \varepsilon = \frac{1}{2} w' w + \frac{1}{2} (\varepsilon - w)' (\varepsilon - w) + w' (\varepsilon - w).$$

Taking expectations under $\hat{\phi}$,

$$\frac{1}{2} \int \varepsilon' \varepsilon \hat{\phi}(\varepsilon) d\varepsilon = \frac{1}{2} w' w + \frac{1}{2} \text{trace}(\Sigma).$$

Combining terms gives

$$(43) \quad \text{ent} = \int (\log \hat{\phi} - \log \phi) \hat{\phi} d\varepsilon = -\frac{1}{2} \log \det \Sigma + \frac{1}{2} w' w + \frac{1}{2} \text{trace}(\Sigma - I).$$

Notice the separate appearances of the mean distortion w and the covariance distortion $\Sigma - I$. We will apply formula (43) to compute a risk-sensitivity operator \mathbb{T} in the next subsection.

A.3. A static valuation problem. In this subsection, we construct a robust estimate of a value function that depends on a random vector that for now we assume is beyond the control of the decision maker. Consider a quadratic value function $V(x) = -\frac{1}{2} x' P x - \rho$ where P is a positive definite symmetric matrix, and $x \sim \mathcal{N}(\bar{x}, \Sigma)$. We shall use the convenient representation $x = \bar{x} + C\varepsilon$, where $CC' = \Sigma$ and $\varepsilon \sim \mathcal{N}(0, I)$. Here $x \in \mathbf{R}^n$, $\varepsilon \in \mathbf{R}^m$, and C is an $n \times m$ matrix.

We want to apply the risk-sensitivity operator \mathbb{T} to the value function $V(x) = -\frac{1}{2} x' P x - \rho$,

$$\mathbb{T}V(\bar{x}) = -\theta \log \int \exp\left(\frac{-V(\bar{x} + C\varepsilon)}{\theta}\right) \phi(\varepsilon) d\varepsilon,$$

where $\phi(\varepsilon) \propto \exp(-\frac{1}{2} \varepsilon' \varepsilon)$ by the assumption that $\phi \sim \mathcal{N}(0, I)$.

Remark 11. For the minimization problem defining $\mathbb{T}V$ to be well posed, we require that θ be sufficiently high that $(I - \theta^{-1} C' P C)$ is nonsingular. The lowest value of θ that satisfies this condition is called the breakdown point.⁴¹

⁴¹See Hansen and Sargent (2008b), chapter 8, for a discussion of the breakdown point and its relation to H_∞ control theory as viewed especially from the frequency domain. See Brock, Durlauf,

To compute TV , we will proceed in two steps.

- *Step 1.* First, we compute $\hat{\phi}(\varepsilon, \bar{x})$. Recall that the associated worst-case likelihood ratio is

$$m(\varepsilon, \bar{x}) \propto \exp\left(\frac{-V(\bar{x} + C\varepsilon)}{\theta}\right),$$

which for the value function $V(x) = -\frac{1}{2}x'Px - \rho$ becomes

$$m(\varepsilon, \bar{x}) \propto \exp\left(\frac{\frac{1}{2}\varepsilon' C' PC\varepsilon + \varepsilon' C' P\bar{x}}{\theta}\right).$$

Then the worst-case density of ε is

$$\begin{aligned} \hat{\phi}(\varepsilon, \bar{x}) &= m(\varepsilon, \bar{x})\phi(\varepsilon) \\ &\propto \exp\left(-\frac{1}{2}\varepsilon'(I - \theta^{-1}C'PC)\varepsilon + \frac{1}{\theta}\varepsilon'(I - \theta^{-1}C'PC)(I - \theta^{-1}C'PC)^{-1}C'P\bar{x}\right). \end{aligned}$$

From the form of this expression, it follows that the worst-case density $\hat{\phi}(\varepsilon, \bar{x})$ is Gaussian with covariance matrix $(I - \theta^{-1}C'PC)^{-1}$ and mean $\theta^{-1}(I - \theta^{-1}C'PC)^{-1}C'P\bar{x} = (\theta I - C'PC)^{-1}C'\bar{x}$.

- *Step 2.* Second, to compute $TV(\bar{x})$, we can use

$$(44) \quad TV(\bar{x}) = \int V(\bar{x} + C\varepsilon)\hat{\phi}(\varepsilon)d\varepsilon + \theta \int m(\varepsilon, \bar{x}) \log m(\varepsilon, \bar{x})\phi(\varepsilon)d\varepsilon$$

while substituting our formulas for the mean and covariance matrix of $\hat{\phi}$ into our formula (43) for the relative entropy of two Gaussian densities. We obtain

$$(45) \quad \begin{aligned} TV(\bar{x}) &= -\frac{1}{2}\bar{x}'\mathcal{D}(P)\bar{x} - \rho - \frac{1}{2}\text{trace}(PC(I - \theta^{-1}C'PC)^{-1}C') + \\ &\frac{\theta}{2}\text{trace}[(I - \theta^{-1}C'PC)^{-1} - I] - \frac{\theta}{2}\log \det(I - \theta^{-1}C'PC)^{-1} \end{aligned}$$

where

$$(46) \quad \mathcal{D}(P) = P + PC(\theta I - C'PC)^{-1}C'P.$$

The matrix $\mathcal{D}(P)$ appearing in the quadratic term in the first line on the right side of (45) emerges from summing contributions coming from (i) evaluating the expected value of the quadratic form $x'Px$ under the worst-case distribution, and (ii) adding in θ times that part of the contribution to entropy $\frac{1}{2}w'w$ in (43) coming from the dependence of the worst-case mean $w = (\theta I - C'PC)^{-1}C'\bar{x}$ on \bar{x} . The term $-\frac{1}{2}\text{trace}(PC(I - \theta^{-1}C'PC)^{-1}C')$ is the usual contribution to the expected value from a quadratic form, but evaluated under the worst-case variance matrix $(I - \theta^{-1}C'PC)^{-1}$. The two terms

and Rondina (2008) for another attack on robust policy design that exploits a frequency domain formulation.

on the second line of (45) are θ times the two contributions from entropy in (43) other than $\frac{1}{2}w'w$.⁴²

Formula (45) simplifies when we note that

$$(I - \theta^{-1}C'PC)^{-1} - I = \theta^{-1}(I - \theta^{-1}C'PC)^{-1}C'PC$$

and that therefore

$$-\frac{1}{2}\text{trace}(PC(I - \theta^{-1}C'PC)^{-1}C') + \frac{\theta}{2}\text{trace}[(I - \theta^{-1}C'PC)^{-1} - I] = 0.$$

So it follows that

$$(47) \quad \text{TV}(\bar{x}) = -\frac{1}{2}\bar{x}'\mathcal{D}(P)\bar{x} - \rho - \frac{\theta}{2}\log \det(I - \theta^{-1}C'PC)^{-1}.$$

It is convenient that with a quadratic objective, linear constraints, and Gaussian random variables, the value function for the risk-sensitivity operator and the associated worst-case distributions can be computed by solving a deterministic programming problem:

Problem 12. *The worst-case mean $v = (\theta I - C'PC)^{-1}C'P\bar{x}$ attains:*

$$\min_v \left\{ -\frac{1}{2}(\bar{x} + Cv)'P(\bar{x} + Cv) + \theta \frac{v'v}{2} \right\}.$$

The minimized value function is $-\frac{1}{2}\bar{x}'\mathcal{D}(P)\bar{x}$ where $\mathcal{D}(P)$ satisfies (51).

A.4. A two-period valuation problem. In this section, we describe a pure valuation problem in which the decision maker does not influence the distribution of random outcomes. We assume the following evolution equation:

$$(48) \quad y^* = Ay + C\varepsilon$$

where y is today's and y^* is next period's value of the state vector, and $\varepsilon \sim \mathcal{N}(0, I)$. There is a value function

$$V(y^*) = -\frac{1}{2}(y^*)'Py^* - \rho.$$

Our risk-sensitive adjustment to the value function is

$$(49) \quad \begin{aligned} \mathbb{T}(V)(y) &= -\theta \log \left[\int \exp \left(\frac{-V[Ay + C\varepsilon]}{\theta} \right) \pi(\varepsilon) d\varepsilon \right] \\ &= \int V(y^*) \hat{\pi} d\varepsilon + \theta \int (\log \hat{\pi} - \log \pi) \hat{\pi} d\varepsilon \end{aligned}$$

⁴²In the special (no-concern about robustness) case that $\theta = +\infty$, we obtain the usual result that

$$\text{TV}(\bar{x}) = EV(\bar{x}) = -\frac{1}{2}\bar{x}'P\bar{x} - \rho - \frac{1}{2}\text{trace}(PCC').$$

To verify this, one shows that the limit of the log det term is the trace term in the second line of (45) as $\theta \rightarrow \infty$. Write the log det as the sum of logs of the corresponding eigenvalues, then take limits and recall the formula expressing the trace as the sum of eigenvalues.

where $\hat{\pi}$ is obtained as the solution to the minimization problem in a multiplier problem. We know that the associated worst-case likelihood ratio satisfies the exponential twisting formula

$$\hat{m}(\varepsilon, y) \propto \exp \left[\frac{1}{2\theta} \varepsilon' C' P C \varepsilon + \frac{1}{\theta} \varepsilon' C' P A y \right].$$

(We have absorbed all non-random terms into the factor of proportionality signified by the \propto sign. This accounts for the dependence of $\hat{m}(\varepsilon, y)$ on y .) When π is a standard normal density, it follows that

$$\pi(\varepsilon) \hat{m}(\varepsilon, y) \propto \exp \left[-\frac{1}{2} \varepsilon' \left(I - \frac{1}{\theta} C' P C \right) \varepsilon + \varepsilon' \left(I - \frac{1}{\theta} C' P C \right) (\theta I - C' P C)^{-1} C' P A y \right],$$

where we choose the factor of proportionality so that the function of ε on the right-hand side integrates to unity. The function on the right side is evidently proportional to a normal density with covariance matrix $(I - \frac{1}{\theta} C' P C)^{-1}$ and mean $(\theta I - C' P C)^{-1} C' P A y$. The covariance matrix of the worst-case distribution is $(I - \frac{1}{\theta} C' P C)^{-1}$ exceeds the covariance matrix I for the original distribution of ε . The altered mean for ε implies that the distorted conditional mean for y^* is $[I + C(\theta I - C' P C)^{-1} C' P] A y$.

Applying (47), the risk-sensitive adjustment to the objective function $-\frac{1}{2}(y^*)' P (y^*) - \rho$ is

$$\begin{aligned} \Gamma(V)(y) &= -\frac{1}{2} (A y)' \mathcal{D}(P) (A y) - \rho \\ (50) \quad &- \frac{\theta}{2} \log \det \left(I - \frac{1}{\theta} C' P C \right)^{-1} \end{aligned}$$

where the operator $\mathcal{D}(P)$ is defined by

$$(51) \quad \mathcal{D}(P) = P + P C (\theta I - C' P C)^{-1} C' P.$$

All of the essential ingredients for evaluating (49) or (50) can be computed by solving a deterministic problem.

Problem 13. *Consider the following deterministic law of motion for the state vector:*

$$y^* = A y + C w$$

where we have replaced the stochastic shock ε in (48) by a deterministic specification error w . Since this is a deterministic evolution equation, covariance matrices do not come into play now, but the matrix C continues to play a key role in designing a robust decision rule. Solve the problem

$$\min_w \left\{ -\frac{1}{2} (A y + C w)' P (A y + C w) + \frac{\theta}{2} w' w \right\}.$$

In this deterministic problem, we penalize the choice of the distortion w using only the contribution to relative entropy (43) that comes from w . The minimizing w is

$$w^* = (\theta I - C'PC)^{-1} C'PAy.$$

This coincides with the mean distortion of the worst-case normal distribution for the stochastic problem. The minimized objective function is

$$-\frac{1}{2}(Ay)'D(P)(Ay),$$

which agrees with the contribution to the stochastic robust adjustment to the value function (50) coming from the quadratic form in Ay . What is missing relative to the stochastic problem is the distorted covariance matrix for the worst-case normal distribution and the constant term in the adjusted value function.

The idea of solving a deterministic problem to generate key parts of the solution of a stochastic problem originated with Jacobson (1973) and underlies much of linear-quadratic-Gaussian robust control theory (for example, see Hansen and Sargent (2008b)). For the purposes of computing and characterizing the decision rules in the linear-quadratic model, we can abstract from covariance distortions and focus exclusively on mean distortions. In the linear-quadratic case, the covariance distortion alters the value function only through the additive constant term $\rho - \frac{1}{2} \log \det(I - \theta^{-1}C'PC)^{-1}$. We can deduce both the covariance matrix distortion and the constant adjustment from formulas that emerge from the purely deterministic problem.

REFERENCES

- Anderson, E. (2005, December). The dynamics of risk-sensitive allocations. *Journal of Economic Theory* 125(2), 93–150.
- Anderson, E., L. Hansen, and T. Sargent (2003). A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *Journal of the European Economic Association* 1(1), 68–123.
- Anderson, E., L. P. Hansen, and T. Sargent (1999, March). Robustness, detection and the price of risk. Mimeo.
- Anderson, E. W., L. P. Hansen, E. R. McGrattan, and T. J. Sargent (1996). Mechanics in forming and estimating dynamic linear economies. In H. M. Amman, D. A. Kendrick, and J. Rust (Eds.), *Handbook of Computational Economics*, pp. 171–252. North Holland.
- Ball, L. (1999). Policy rules for open economies. In J. Taylor (Ed.), *Monetary Policy Rules*, pp. 127–144. University of Chicago Press.
- Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance* LIX(4), 1481–1509.

- Barillas, F., L. P. Hansen, and T. J. Sargent (2009, November). Doubts or variability? *Journal of Economic Theory* 144(6), 2388–2418.
- Barsky, R. T., T. Juster, M. Kimball, and M. Shapiro (1997). Preference parameters and behavior heterogeneity: An experimental approach in the health and retirement survey. *Quarterly Journal of Economics* 112, 537 – 580.
- Basar, T. and P. Bernhard (1995). *H[∞]-Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach*. Birkhauser.
- Blackwell, D. A. and M. A. Girshick (1954). *Theory of Games and Statistical Decisions*. London, England: John Wiley and Sons.
- Brainard, W. (1967). Uncertainty and the effectiveness of policy. *American Economic Review* 57, 411–425.
- Brock, W. A., S. N. Durlauf, and G. Rondina (2008, May). Frequency-specific effects of stabilization policies. *American Economic Review* 98(2), 241–45.
- Brock, W. A. and L. Mirman (1972). Optimal economic growth and uncertainty: the discounted case. *Journal of Economic Theory* 4(3), 479–513.
- Cagetti, M., L. P. Hansen, T. J. Sargent, and N. Williams (2002, March). Robustness and pricing with uncertain growth. *Review of Financial Studies* 15(2), 363–404.
- Camerer, C. (1995). Individual decision making. In J. H. Kagel and A. Roth (Eds.), *Handbook of Experimental Economics*, pp. 588–673. Princeton University Press.
- Camerer, C. (1999). Ambiguity-aversion and non-additive probability: Experimental evidence, models and applications. In L. Luini (Ed.), *Uncertain Decisions: Bridging Theory and Experiments*, pp. 53–80. Kluwer Academic Publishers.
- Carboni, G. and M. Ellison (2009, September). The great inflation and the greenbook. *Journal of Monetary Economics* 56(6), 831–841.
- Chen, Z. and L. Epstein (1998, August). Ambiguity, risk, and asset returns in continuous time. Mimeo.
- Chernoff, H. (1952). A measure of asymptotic efficiency for tests of a hypothesis based on sums of observations. *Annals of Mathematical Statistics* 23, 493–507.
- Christiano, L. J. and T. J. Fitzgerald (1998). The business cycle: Its still a puzzle. *Federal Reserve Bank of Chicago Economic Perspectives IV*, 56–83.
- Cochrane, J. H. and L. P. Hansen (1992). Asset pricing explorations for macroeconomics. *NBER Macroeconomics Annual* 7, 115–169.
- Cogley, T., R. Colacito, L. P. Hansen, and T. J. Sargent (2008, December). Robustness and u.s. monetary policy experimentation. *Journal of Money, Credit and Banking* 40(8), 1599–1623.
- Cogley, T., R. Colacito, and T. J. Sargent (2007, 02). Benefits from u.s. monetary policy experimentation in the days of samuelson and solow and lucas. *Journal of Money, Credit and Banking* 39(s1), 67–99.

- Ellsberg, D. (1961). Risk, ambiguity and the savage axioms. *Quarterly Journal of Economics* 75, 643–669.
- Epstein, L. G. and A. Melino (1995). A revealed preference analysis of asset pricing under recursive utility. *Review of Economic Studies* 62, 597–618.
- Epstein, L. G. and T. Wang (1994). Intertemporal asset pricing under knightian uncertainty. *Econometrica* 62(3), 283–322.
- Ferguson, T. S. (1967). *Mathematical Statistics: A Decision Theoretic Approach*. New York, NY: Academic Press.
- Fleming, W. and P. Souganidis (1989). On the existence of value functions of two-player, zero-sum stochastic differential games. *Indiana Univ. Math. J.* 38, 293–314.
- Friedman, M. (1953). The effects of a full-employment policy on economic stability: A formal analysis. In M. Friedman (Ed.), *Essays in Positive Economics*. University of Chicago Press.
- Friedman, M. (1956). *A Theory of the Consumption Function*. Princeton University Press.
- Friedman, M. (1959). *A Program for Monetary Stability*. Fordham University Press.
- Friedman, M. and L. J. Savage (1948). The utility analysis of choices involving risk. *Journal of Political Economy* 56, 279–304.
- Gilboa and Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18, 141–153.
- Hansen, L. P. and R. Jagannathan (1991). Implications of security market data. *Journal of Political Economy* 99, 225–261.
- Hansen, L. P. and T. Sargent (1995a, May). Discounted linear exponential quadratic gaussian control. *IEEE Transactions on Automatic Control* 40(5), 968–971.
- Hansen, L. P., T. Sargent, and T. Tallarini (1999). Robust permanent income and pricing. *Review of Economic Studies* 66, 873–907.
- Hansen, L. P. and T. J. Sargent (1993). Seasonality and approximation errors in rational expectations models. *Journal of Econometrics* 55, 21–55.
- Hansen, L. P. and T. J. Sargent (1995b). Discounted linear exponential quadratic gaussian control. *IEEE Transactions on Automatic Control* 40, 968–971.
- Hansen, L. P. and T. J. Sargent (2001, May). Robust control and model uncertainty. *American Economic Review* 91(2), 60–66.
- Hansen, L. P. and T. J. Sargent (2007, September). Recursive robust estimation and control without commitment. *Journal of Economic Theory* 136(1), 1–27.
- Hansen, L. P. and T. J. Sargent (2008a). Fragile beliefs and the price of uncertainty. University of Chicago and New York University.
- Hansen, L. P. and T. J. Sargent (2008b). *Robustness*. Princeton, New Jersey: Princeton University Press.

- Hansen, L. P., T. J. Sargent, G. A. Turmuhambetova, and N. Williams (2006, March). Robust control, min-max expected utility, and model misspecification. *Journal of Economic Theory* 128, 45–90.
- Hansen, L. P., T. J. Sargent, and N. E. Wang (2002). Robust permanent income and pricing with filtering. *Macroeconomic Dynamics* 6, 40–84.
- Harlevy, Y. (2007). Ellsberg revisited: An experimental study. *Econometrica* 75(2), 503–536.
- Jacobson, D. H. (1973). Optimal linear systems with exponential performance criteria and their relation to differential games. *IEEE Transactions on Automatic Control* 18, 124–131.
- Jovanovic, B. (1979). Job matching and the theory of turnover. *Journal of Political Economy* 87(5), 972–990.
- Jovanovic, B. and Y. Nyarko (1996). Learning by doing and the choice of technology. *Econometrica* 64(6), 1299–1310.
- Karantounias, A. G. (2009). Managing expectations and fiscal policy. Working paper, Federal Reserve Bank of Atlanta. Lars Peter Hansen and Thomas J. Sargent.
- Kasa, K. (1999, March). Model uncertainty, robust policies, and the value of commitment. Mimeo.
- Keynes, J. M. (1936). *The General Theory of Employment, Interest, and Money*. Macmillan.
- Kimball, M. (1990). Precautionary saving in the small and in the large. *Econometrica* 58, 53–73.
- Knight, F. H. (1921). *Risk, Uncertainty and Profit*. Houghton Mifflin Company.
- Kreps, D. M. (1998). Anticipated utility and dynamic choice. In D. P. Jacobs, E. Kalai, and M. I. Kamien (Eds.), *Frontiers of research in economic theory: the Nancy L. Schwartz Memorial Lectures, 1983-1997*. Cambridge University Press.
- Leland, H. (1968). Savings and uncertainty: The precautionary demand for savings. *Quarterly Journal of Economics* 82, 465–473.
- Levin, A. T. and J. C. Williams (2003). Robust monetary policy with competing reference models. *Journal of Monetary Economics* 50(5), 945–975.
- Ljungqvist, L. and T. J. Sargent (2004). *Recursive Macroeconomic Theory, Second edition*. Cambridge, Massachusetts: MIT Press.
- Lucas, R. E. (1976). Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy: The Phillips Curve and Labor Markets* 1, 19–46.
- Marcet, A. and J. P. Nicolini (2003, December). Recurrent hyperinflations and learning. *American Economic Review* 93(5), 1476–1498.
- Miller, B. L. (1974). Optimal consumption with a stochastic income stream. *Econometrica* 42, 253–266.

- Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica* 29, 315–335.
- Onatski, A. and J. H. Stock (1999). Robust monetary policy under model uncertainty in a small model of the u.s. economy. Mimeo.
- Onatski, A. and N. Williams (2003). Modeling model uncertainty. *Journal of the European Economic Association* 1(5), 1087–1122.
- Orlik, A. and I. Presno (2009). On credible monetary policies with model uncertainty. Mimeo, New York University.
- Petersen, I. R., M. R. James, and P. Dupuis (2000). Minimax optimal control of stochastic uncertain systems with relative entropy constraints. *IEEE Transactions on Automatic Control* 45, 398–412.
- Rabin, M. (1999, January). Risk aversion and expected-utility theory: A calibration theorem. Mimeo.
- Sargent, T., N. Williams, and T. Zha (2006, September). Shocks and government beliefs: The rise and fall of american inflation. *American Economic Review* 96(4), 1193–1224.
- Sargent, T., N. Williams, and T. Zha (2009, 04). The conquest of south american inflation. *Journal of Political Economy* 117(2), 211–256.
- Sargent, T. J. (1981). Interpreting economic time series. *Journal of Political Economy* 89(2), 213–248.
- Sargent, T. J. (1999a). Comment. In J. Taylor (Ed.), *Monetary Policy Rules*, pp. 144–154. University of Chicago Press.
- Sargent, T. J. (1999b). *The Conquest of American Inflation*. Princeton University Press.
- Sargent, T. J. and N. Wallace (1975). Rational expectations, the optimal monetary instrument, and the optimal supply of money. *Journal of Political Economy* 83, 241–254.
- Savage, L. J. (1954). *The Foundations of Statistics*. John Wiley and Sons.
- Segal, U. and A. Spivak (1990). First order versus second order risk aversion. *Journal of Economic Theory* 51, 111 – 125.
- Sims, C. A. (1972). Approximate prior restrictions in distributed lag estimation. *Journal of the American Statistical Association* 67, 169–175.
- Sims, C. A. (1974, September). Seasonality in regression. *Journal of the American Statistical Association* 69(347), 618–626.
- Sims, C. A. (1993). Rational expectations modeling with seasonally adjusted data. *Journal of Econometrics* 55, 9–19.
- Tallarini, T. D. (2000, June). Risk sensitive business cycles. *Journal of Monetary Economics* 45(3), 507–532.
- Taylor, J. B. and J. C. Williams (2009). Simple and robust rules for monetary policy. draft for KDME Conference October 2009.

- von Neumann, J. and O. Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton University Press.
- Wallis, K. F. (1974, March). Seasonal adjustment and relations between variables. *Journal of the American Statistical Association* 69(345), 18–31.
- Whiteman, C. H. (1986). Analytical policy design under rational expectations. *Econometrica* 54, 1387–1405.
- Whittle, P. (1980). *Risk-Sensitive Optimal Control*. John Wiley and Sons.
- Wieland, V. (1996). Monetary policy, parameter uncertainty, and optimal learning. Mimeo: Board of Governors of the Federal Reserve System.
- Woodford, M. (2010, March). Robustly optimal monetary policy with near-rational expectations. *American Economic Review* 100(1), 274–303.
- Zeldes, S. P. (1989). Optimal consumption with stochastic income: Deviation from certainty equivalence. *Quarterly Journal of Economics* 104, 275–298.
- Zhou, K., J. C. Doyle, and K. Glover (1996). *Robust and Optimal Control*. Prentice-Hall.

LARS PETER HANSEN, DEPARTMENT OF ECONOMICS, UNIVERSITY OF CHICAGO, CHICAGO, ILLINOIS

E-mail address: l-hansen@uchicago.edu

THOMAS J. SARGENT, DEPARTMENT OF ECONOMICS, NEW YORK UNIVERSITY AND HOOVER INSTITUTION, STANFORD UNIVERSITY, STANFORD, CALIFORNIA

E-mail address: ts43@nyu.edu