

WANTING ROBUSTNESS IN MACROECONOMICS

LARS PETER HANSEN AND THOMAS J. SARGENT

1. INTRODUCTION

1.1. **Foundations.** von Neumann and Morgenstern (1944) and Savage (1954) created mathematical foundations that applied economists have used to construct quantitative dynamic models for policy making and empirical analyses. The mathematical foundations give modern dynamic models internal coherence and sharp empirical predictions. However, those foundations should invite researchers to confront the unsettling fact that their models are approximations. That would expose logical problems that until recently have been swept under the rug, but that still call for repair.

A model is a probability distribution over a sequence. Applied dynamic economists readily accept that their models are approximations¹ because applied models must be *tractable*, that is, feasible to solve, estimate, and simulate. With tractability comes misspecification. Model misspecification is unavoidable in applied economic research. Admitting model misspecification raises especially interesting problems for rational expectations models because the rational expectations assumption excludes model misspecification.

The rational expectations hypothesis imposes a communism of models: the people being modeled know the model. This makes the economic analyst, the policy maker, and the agents being modeled all share the same model, *i.e.*, the same probability distribution over sequences of outcomes. Imposition of a common model removes economic agents' models as objects that require separate specification. 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 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 little work studying how

Date: June 17, 2000.

We thank François Velde and Neng Wang for insightful comments on an earlier draft.

¹Sometimes they express this by saying that their models are abstractions or idealizations. Other times they convey it by focusing their model only on 'stylized facts'.

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

Thus, the idea that models are approximations conflicts both with the von Neumann-Morgenstern-Savage foundations and the supplementary equilibrium concept of rational expectations. 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 the passage from Keynes (1936) which they cite:

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, which might explain the sudden fluctuation in opinion that Keynes mentions. Our approach provides no such analysis of the sudden fluctuation Keynes wrote about. However, our analysis of detection error probabilities gives a way to think about the shroud of uncertainty around the distribution of future yields.

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

In *Risk, Uncertainty and Profit*, Frank Knight (1921) envisioned entrepreneurs who confront a form of uncertainty not captured by a probability model as they search for profits.² He distinguished between risk and uncertainty, and reserved the term risk for ventures for which probabilities are known. 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) devised an axiomatic treatment of decision-making in which actions could be justified as maximizing expected utility with the

²See Epstein and Wang (1994) for a discussion containing many of the ideas summarized here.

assignment of subjective probabilities. Savage's work extended the earlier justification of expected utility by von Neumann and Morgenstern (1944) using prespecified probabilities. That is, Savage's axioms also justified (subjective) probability assignments. Even when accurate probabilities, such as the fifty-fifty put on the sides of a fair coin, were not available, decision makers conforming to Savage's axioms behave as if they formed 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. Moreover, it provides a direct, but possibly uninteresting, recipe for confronting *misspecified* models. 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 replace the known density f by a family of densities $g(y|x, d, \alpha)$ indexed by parameters α giving a family of potential models. By averaging over the array of candidate models using a prior (subjective) distribution, say π , we can form a 'hyper model' taken to be 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.

2.2. Savage and rational expectations. Rational expectations theory withdrew some freedom from Savage's decision theory by imposing consistency 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.³ 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. By undermining 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)

³For example, see Sargent (1981).

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. 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 position. This example, called the *Ellsberg paradox*, challenges the appropriateness of the full array of Savage axioms.⁴

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 two gambles for 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 for this mixture. Under *aversion to uncertainty*, mixing with known probabilities can only improve the welfare of the decision-maker. Thus Gilboa and Schmeidler asked that the decision-maker at least weakly prefer the mixture of gambles to either of the original gambles.

The resulting generalized decision theory is supported by 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

⁴In 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.




FIGURE 1. The Ellsberg Urn

on the decision being considered and is designed as part of the decision making process 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.

The outcome of the Gilboa-Schmeidler axioms is a decision theory that can be formalized as a two-player game. For every action of one player, a second 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. 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 decision-maker who selects beliefs in a pessimistic fashion. Restrictions on beliefs must come from outside.

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 it⁵ 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:

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.⁶ 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 will draw our analysis from a line of research in control theory. Robust control theorists challenged modern control theory because it ignored model-approximation error in designing policy rules. They feared that their models misspecified the dynamic responses of target variables to controls. They added an error process to their base models, and sought decision rules that would work well across a set of such disturbance 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.

⁵Unlike Lucas (1976) and Sargent and Wallace (1975).

⁶However, 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 vector of independent and identically normally distributed *shocks* with mean zero and covariance matrix given by I . The target vector is used to define preferences via:

$$(3) \quad -\frac{1}{2} \sum_{t=0}^{\infty} \beta^t E|z_t|^2$$

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 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:

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

subject to

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

where w^* has mean zero and covariance matrix I .⁷ Here * superscripts denote next-period values. This is a fixed-point problem because the same positive semidefinite matrix Ω and scalar ω occur on both the right and left sides. The value function for this decision problem is the solution to the fixed point problem: $V(x) = -\frac{1}{2}x'\Omega x - \omega$.

The solution of the ordinary linear quadratic optimization problem has a special property called *certainty equivalence*. Certainty equivalence means that the decision rule F is independent of the ‘noise statistics’. We state this formally in

Claim 2. (*Certainty Equivalence Principle*)

For the linear-quadratic control problem, the matrix Ω^o and the optimal control law F^o do not depend on the noise statistics embedded in C . Thus the optimal control law does not depend on the matrix C .

⁷There 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. In seeking to improve controls, they focus on the specification of the shock process.

Can the 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 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. As we know from the dynamic programming formulation of the single-agent decision problem, it is easiest 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-agent 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 the 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, as in Gilboa and Schmeidler (1989), a second-agent has entered the analysis. This agent aims to *minimize* the objective, but in doing so is penalized by the 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).

The fictitious second agent puts *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 magnitude of the pessimistic distortion is restrained by the robustness parameter θ . 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.

The solution of this game includes a worst case v , which is a function of the state x and the control u , as we will see below. Since v is allowed to depend on the current period state and control, the resulting decision problem allows for misspecification of the dynamic evolution of x . Moreover, the matrix C is now an impact matrix for misspecification, and contributes to the *robust* solution to the control problem. We summarize the dependence of F on C in the following, which is to be contrasted with claim 2:

Claim 4. (*Breaking Certainty Equivalence*)

For $\theta < +\infty$, the robust control $u = -Fx$ that solves game (3) depends on the noise statistics as intermediated through 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.

3.3. Gilboa and Schmeidler again. To relate formulation (3) to that of Gilboa and Schmeidler (1989), we look at a hybrid specification in which both the shock w^* and the specification error v are present. The idea is to replace the i.i.d. shock w^* with $w^* + v$, where v is permitted to be a function of current and past values of the state vector. Thus, v can be viewed as the mean of the shock $v + w^*$ conditioned on the current state x . Alternative v 's index the *multiple priors* in the date t evolution equation. Instead of restraining the v 's to reside in some prespecified set, for convenience we penalize their magnitude directly in the objective function.⁸

This formulation of robustness can capture a variety of model misspecifications. Consider, for instance, misspecification of the matrices A and B . Suppose instead the evolution is:

$$x^* = A^*x + B^*u + Cw^*.$$

Then we can write:

$$x^* = Ax + Bu + C(v + w^*),$$

provided that we can solve the equation:

$$Cv = (A^* - A)x + (B^* - B)u$$

for some v . The matrix C limits the form of the misspecification of A and B , for reasons that we will discuss later.

⁸As discussed in Hansen, Sargent, and Tallarini (1999) and Hansen and Sargent (2000), we can think of the robustness parameter θ as a Lagrange multiplier on an intertemporal specification-error constraint.

More general forms of misspecified dynamics can be entertained when solving the above game, including nonlinearities, higher-order dynamics and time variation in the parameters. As a consequence this formulation of robustness is called *unstructured*.⁹ Nevertheless, the constrained (by θ) worst-case model has the form of the original model with an altered A and B . Associated with the solution to the two-player game is a worst-case choice of v . This worst case corresponds to a particular (A^*, B^*) 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:

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

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:

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

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

4. CALIBRATING THE TASTE FOR ROBUSTNESS

Our model of a robust decision-maker is formalized as a two-player game with a fictitious new player who aims to deceive. This new player, 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 promote skepticism about the use of minimization rather than the averaging advocated in Bayesian decision theory.

Whether the formulation in terms of the game looks silly or plausible depends on how the fictitious new player is disciplined. While an undisciplined malevolent player can *wreak havoc*, a tightly constrained one will have less consequence. 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. There is some support for making conservative adjustments in experimental evidence (see Camerer (1995) for a discussion) and in the axiomatic treatment of Gilboa and Schmeidler (1989). Neither of these sources answer the quantitative question of

⁹See Stock and Onatski (1999) for an example of robust decision analysis with structured uncertainty.

how big the adjustment should be in applied work in economic dynamics. Here the theory of statistical discrimination can help.

We have parameterized a taste for robustness in terms of a single free parameter, θ . The parameter θ indexes a family of robust decision problems. We want to know which members of this family 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 match the model that generates the actual time series evolution. A defense of the rational expectations equilibrium is that discrepancies between models could have been detected from sufficient historical data. In this section, we use this same approach to think about reasonable preferences for robustness. Given historical observations on the state vector, we use the Bayesian 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. 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 of the state vector x_t , suppose that we want to determine the evolution equation for the state vector. Let $u = -F^*x$ denote the solution to the robust control problem. One possible description of the time series is

$$(7) \quad x_{t+1} = (A - BF^*)x_t + Cw_{t+1}.$$

In this case, concerns about model misspecification are just in the head of the decision-maker: the original model is actually correctly specified, so that $v^* \equiv 0$. Here the approximating model actually generates the data.

An alternative evolution equation is the one associated with the solution to the two-player game. In this case, write (6) as

$$v^* = -K^*x$$

where

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

Then

$$(8) \quad x_{t+1} = (A - BF^* - CK^*)x_t + Cw_{t+1}.$$

This evolution takes the constrained worst case model as the actual law of motion of the state vector, under the robust decision rule and the worst-case shock process

that the decision maker plans against.¹⁰ Since the choice of v by the new agent is not meant to be a prediction, only a conservative adjustment, this evolution is not the decision maker's guess about the most likely model. The decision maker believes that the evolution is actually some *unknown* law

$$(9) \quad x_{t+1} = (A - BF^*)x_t + C(v_t + w_{t+1}),$$

where the v_t 's are of no larger magnitude than the v_t^* 's in the sense that

$$E_0 \sum_{t=0}^{\infty} \beta^t v_t^{*'} v_t^* \geq E_0 \sum_{t=0}^{\infty} \beta^t v_t' v_t,$$

where E_0 is evaluated with respect to the corresponding model. The actual v_t 's could be those associated with a much more complicated form of model misspecification than the solution to the two-player game. Nevertheless, the two evolution equations (7) and (8) provide a convenient laboratory for gauging 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) . Suppose another is indexed by $\{v_t\}$ in (9). When thinking about statistical detection, it is imperative that we include some actual randomness, in contrast to many formulations of robust control theory. That 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. Thus we include a stochastic shock w_{t+1} to disguise the model misspecification v_t .¹¹

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 log likelihood ratio is

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

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

¹⁰It is the decision rule from the Markov perfect equilibrium of the dynamic game.

¹¹Here, for pedagogical convenience we explore only a special stochastic departure from the approximating model. As emphasized by Anderson, Hansen, and Sargent (1999), 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.

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, and 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 on the state x_t , so that to simulate them requires simulating a particular law of motion (9).

Use of this measure 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 counterpart. 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 normal model. 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 is absent in many formulations of robust control theory.

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

We believe that the model detection problem confronted by a decision-maker is actually more complicated than the pair-wise statistical discrimination problem we just described. A decision-maker will instead 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 process $\{v_t^*\}$ and consider

only those values of θ for which the $\{v_t^*\}$ 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^* 's than of the θ 's that underlie them. This suggests solving robust control problems for an array of θ 's and exploring the resulting v_t^* 's.

4.3.1. *Detection probabilities: an example.* Here is how we construct detection error probabilities in practice. Consider two alternative models. Model A is the approximating model (with $v_t \equiv 0$ in (9)). Let model B be the worst case model associated with $v_t = K^*(\theta)x_t$, for a particular positive θ . Consider a fixed sample of T observations on x_t . Let L_{ij} be the likelihood of that sample for model j assuming that model i generates the data. Define the likelihood ratio

$$r_i = \log \frac{L_{ii}}{L_{ij}}$$

where $j \neq i$ and $i = A, B$. 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 mistakes. First, assume that model A generates the data and calculate

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

Next, assume that model B generates the data and calculate

$$p_B = \text{Prob}(\text{mistake}|B) = \text{freq}(r_B \leq 0).$$

Under a prior that 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 (2000) for more details about computing detection error probabilities in various linear models.

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.¹²

¹²See Sargent (1999a) for further discussion of Ball's model from the perspective of robust decision theory. See Hansen and Sargent (2000) for how to treat robustness in 'forward looking' models.

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

$$(10) \quad y_t = -\beta r_{t-1} - \delta e_{t-1} + \epsilon_t$$

$$(11) \quad \pi_t = \pi_{t-1} + \alpha y_{t-1} - \gamma(e_{t-1} - e_{t-2}) + \eta_t$$

$$(12) \quad e_t = \theta r_t + \nu_t,$$

where y is the log of real output, r is the real interest rate, e is the log 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 the monetary authority wants to maximize

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

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

Ball motivates (10) as an open-economy IS curve and (11) as an open-economy Phillips curve; he uses (12) to capture effects of the interest rate on the exchange rate. Ball set the parameters $\gamma, \theta, \beta, \delta$ at the values .2, 2, .6, .2. Following Ball, we set the innovation shock standard deviations equal to 1, 1, $\sqrt{2}$.

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 want 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'll 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. Our formulation treats misspecification of all of the state-evolution equations symmetrically. All types of misspecification that can be disguised by the shock vector w_{t+1} are admitted. 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




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

in some aspects of the model than others. More generally, multiple 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.

This use of the theory of statistical discrimination helps to excuse a decision not to model active learning about model misspecification. Next we explore ways of incorporating learning.

5. LEARNING

The robust control theoretic 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 that were available in stage 1 and that are likely to be available until a long time has passed in stage 2.

This section considers approaches to model ambiguity coming from literatures on adaptation. They do not temporally separate learning from control. 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 A and B as being random and specify a prior distribution for this randomness. One possibility is that there is only some initial randomness. 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).

Alternatively, we might envision a stochastic process of the (A, B) 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) model. Specification of the problem now becomes more demanding, as the decision-maker is compelled to take a stand on the stochastic evolution of the matrices (A, B) . The solutions are also 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 in the next section.

5.2. Adaptive Models. Another approach to model uncertainty comes from distinct literatures on adaptive control and vector autoregressions with random coefficients.¹³ 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

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

and that the coefficients are described by

$$(13) \quad \begin{bmatrix} \text{col}(A_{t+1}) \\ \text{col}(B_{t+1}) \end{bmatrix} = \begin{bmatrix} \text{col}(A_t) \\ \text{col}(B_t) \end{bmatrix} + \begin{bmatrix} \eta_{A,t+1} \\ \eta_{B,t+1} \end{bmatrix}$$

¹³See Kreps (1998) and Sargent (1999b) for related accounts of this approach. See Marcat and Nicolini (1997) for an empirical application.

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 setting 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

$$(14) \quad x^* = \hat{A}_t x + \hat{B}_t u + C w^*.$$

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 that assumption that A, B 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 looking into the future while making decisions today. He is of two minds: when determining the control $u_t = -F x_t$ 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 believed (13). It is often excused because it is much simpler than a Bayesian analysis.

5.3. 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.¹⁴ 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 3 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 $d_t = d_t^1 + d_t^2$.

Figure 4 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 5 uses Bayesian updating (Kalman filtering) forms 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 5 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 in technological growth.¹⁵ For example, investors may not have perfect information about the timing of the productivity slowdown that probably occurred in the

¹⁴See Jovanovic (1979) Jovanovic and Nyarko (1996) for examples of this idea.

¹⁵It 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 (2000) 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.

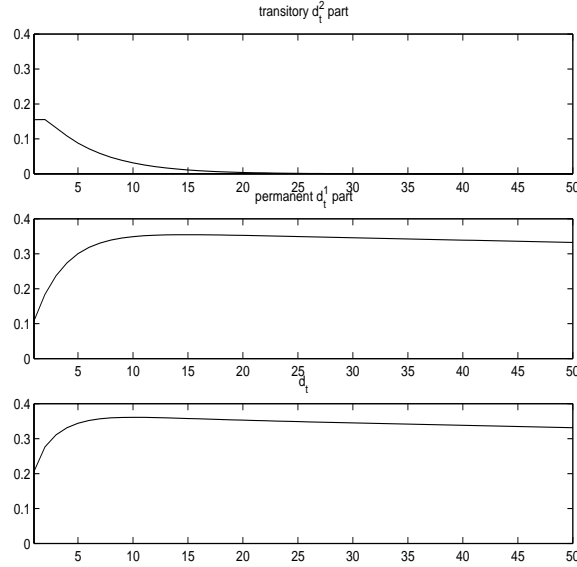


FIGURE 3. Impulse responses for two components of endowment process and their sum in Hansen, Sargent, and Tallarini's model. The top panel is the impulse response of the transitory component d^2 to an innovation in d^2 ; the middle panel, the impulse response of the permanent component d^1 to its innovation; the bottom panel is the impulse response of the sum $d_t = d_t^1 + d_t^2$ to its own innovation.




FIGURE 4. Actual permanent and transitory components of endowment process from Hansen, Sargent, Tallarini (1999) model.




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

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 6 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.

5.4. 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 we described previously. 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 predicted state. In a stochastic counterpart to a steady state, the predicted state evolves according to:

$$(15) \quad \hat{x}^* = A\hat{x} + Bu + G_x \hat{w}^*$$

$$(16) \quad y^* = SA\hat{x} + SBu + G_y \hat{w}^*$$

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

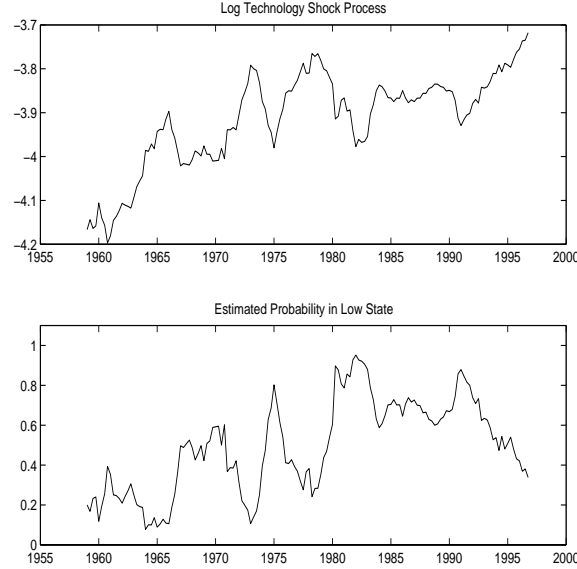


FIGURE 6. Top panel: the growth rate of the Solow residual, a measure of the rate of technological growth. Bottom panel: the probability that growth rate of the Solow residual is in the low-growth state.

next-period's data y^* via the formula

$$(17) \quad \hat{w} = (G_y)^{-1}(y^* - SA\hat{x} - SBu).$$

However, the original w^* cannot generally be recovered. 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 .¹⁶

In many decision problems confronted by macroeconomists, the target depends only on the observable component of the state, and thus:¹⁷

$$(18) \quad z = H\hat{x} + Ju,$$

5.5. 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 (15) and target equation (18). These replace the original state evolution equation (1) and target equation (2). The G_x matrix replaces the C matrix, but because of

¹⁶In fact, the matrices G_x and G_y are not unique but the so-called gain matrix $K = G_x(G_y)^{-1}$ is.

¹⁷A more general problem in which z depends directly on hidden components of the state vector can also be handled.

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 (predicted) state \hat{x} generated recursively by the Kalman filter.

5.6. 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 (2000) 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 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 (2000), Hansen, Sargent, and Wang (2000) and Cagetti, Hansen, Sargent, and Williams (2000) for an elaboration.

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

To compute a robust control law, we solve robustness game 3 but with the information or predicted state \hat{x} replacing the original state x . Since we perturb evolution equation (15) instead of (1), we substitute the matrix G_x for C when solving the robust control problem. Since the equilibrium of our earlier robustness 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.¹⁸ The separation of prediction 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 (15), (16) 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.

¹⁸Although 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 $G_x(G_x)'$.

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:

$$(19) \quad x^* = A\hat{x} + Bu + \check{G}\check{w}^*,$$

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

To study robustness, we disguise the model misspecification by the shock \check{w}^* . Notice that the dimension of \check{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 (19) 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 robustness 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 robustness 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 predictor \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 3.¹⁹ The fact that the decision-makers explore

¹⁹Game 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.

the evolution of x^* instead of the information state \hat{x}^* gives robustness more potency in game 8 than 7.²⁰

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.

5.7. Adaptation versus robustness. 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 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 7 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 7 plots the value $C = -E(\pi^2 + y^2)$ attained by three rules under the alternative data generating model associated with 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.

²⁰As emphasized in Hansen, Sargent, and Wang (2000) holding θ fixed across games is different than holding detection errors probabilities fixed.




FIGURE 7. Value of $C = -E(\pi^2 + y^2)$ 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.

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 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:

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

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:

$$(20) \quad z_t = (H - JF)(A - BF)^t C w_0.$$

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:

$$(21) \quad -\frac{1}{2}(w_0)' \sum_{t=0}^{\infty} \beta^t \mathcal{I}(F, t - 1)' \mathcal{I}(F, t - 1) w_0.$$

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 $\mathcal{I}(F, 0), \sqrt{\beta}\mathcal{I}(F, 1), \dots, \sqrt{\beta}^t \mathcal{I}(F, t), \dots$ small in magnitude. Thus (21) induces no preferences over specific patterns of the impulse response sequence, only about the overall magnitude of the sequence.

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 x_0 , the solution of the problem of choosing a fixed F to maximize (21) 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.*

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 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 initialization. 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, the choice of the v_t sequence is penalized using robustness parameter θ . Thus the non recursive objective for the

dynamic game is:

$$(22) \quad - \sum_{t=0}^{\infty} \beta^t \{ |z_t| - \theta |v_t|^2 \}.$$

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. Thus we now have a (non recursive) dynamic game in which a new fictitious agent chooses a v_t sequence to minimize (22). To construct a robust law, the original decision-maker then maximizes (22) by choice of F . This non recursive representation of the game yields to the Fourier transform techniques studied by Whiteman (1986), Kasa (1999), and Christiano and Fitzgerald (1998). See Hansen and Sargent (2000) for a formal development. Moreover, this non recursive game has the same solution as the recursive game 3.

Before describing some details, it is easy to describe how this second agent will behave. He will find 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 in 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 leaves himself most vulnerable. Here the mathematical tool of Fourier transforms allows us to summarize the impulse response function in the frequency domain. Imagine using sine/cosine specifications of the components of the sequence v_t 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 allow the minimizing agent to put specific temporal patterns in the v_t 's. For such cosine wave shocks, there is no simple uncoupling like that which occurs with random shocks or impulses. It is necessary to view the composite contribution of entire v_t ' sequence, including its temporal pattern. The impulse response sequence summarizes how future targets respond to a current period v_t ; the 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 such temporally dependent v_t sequences, the maximizing agent should then care about the temporal pattern of the impulse response sequence, not just its overall magnitude.²¹ 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 expose model misspecification. 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. A trade-off across frequencies now emerges. The robustness

²¹Thus, it 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.

parameter θ balances a tension between asking that impulse responses be of small magnitude and also that they be insensitive to model misspecification.

6.3. Some frequency domain details. To investigate the preceding 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:

$$\mathcal{FT}(v)(\theta) = \sum_{t=0}^{\infty} \beta^{t/2} v_t \exp(i\theta t)$$

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 (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 us with a convenient toolkit for thinking formally about misspecification using frequency decompositions.

It may appear troubling that the frequency components are complex. However, by combining the contribution of frequency ω with frequency $-\omega$ we obtain sequence 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.²² This calculation results in the formula:

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

²²That 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.

where, in particular,

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

The matrix function h is the Fourier transform of the sequence of impulse responses. 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.

Objective (22) 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 maximize 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 (2000), this is equivalent to ranking control laws F using the frequency-based *entropy* criterion:

$$(23) \quad \text{entropy} = -\frac{1}{2\pi} \int_{-\pi}^{\pi} \log \det[\theta I - h(\omega)'h(-\omega)] d\omega.$$

See Hansen and Sargent (2000) 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, and to lower 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 game suffers a *breakdown*. The fictitious malevolent player can inflict sufficient damage that the objective function remains at $-\infty$ independent of the control law. The critical value of θ can be found by solving:

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

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 θ^* depends on the choice of control law F . One (somewhat extreme) version of robust control theory called H^∞ control theory would have a decision-maker select the 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 are often 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 suited 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.

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.

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 defend filtering 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 robustness formulation for this defense. If an econometrician suspects that a model is better specified at some frequencies than others, this can be operationalized by allowing the fictitious 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 analagous two-player game. To blunt the role of this mischief-making, the econometrician will design a filter to eliminate those frequencies from estimation.

Sims's analysis provides a way to think about both seasonal adjustment and trend removal. Both can be regarded as procedures that remove high powered frequency

components with the aim of focusing empirical analysis on frequencies where a model is better specified. For the case of seasonal adjustment, Sims (1993) and Hansen and Sargent (1993) in effect describe a situation in which, through the cross-equation restrictions of rational expectations models, b is a function of preference and technological parameters that the econometrician wants to estimate. Hansen and Sargent (1993) provide examples of models whose poorer specifications at the seasonal components can inspire a preference for estimators constructed using 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 Sims's analysis of data filtering, an 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 how the v_t process feeds back on the state depends on his own 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 preferences for robustness affect frequency domain representations of impulse




FIGURE 8. Frequency decompositions of $[\text{trace } h(\omega)'h(-\omega)]$ for objective function of Ball's model under three decision rules; $\beta = 1$ on left panel, discount factor $\beta = .9$ on right panel.

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

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

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 8 displays frequency decompositions of $[\text{trace } h(\omega)'h(-\omega)]$ for robust rules with $\beta = 1$ and $\beta = .9$. Figure 8 shows frequency domain decompositions of a government's objective function for three alternative policy rules labelled $\theta = \infty, \theta = 10, \theta = 5$. The parameter θ measures a preference for robustness, with $\theta = +\infty$ corresponding to the no preference for robustness, and lower values of θ representing a concern for misspecification. Of the three rules whose transfer functions are depicted in figure Figure 8, Ball's rule ($\theta = \infty$) is the best under the approximating model because the area under the curve is smallest.

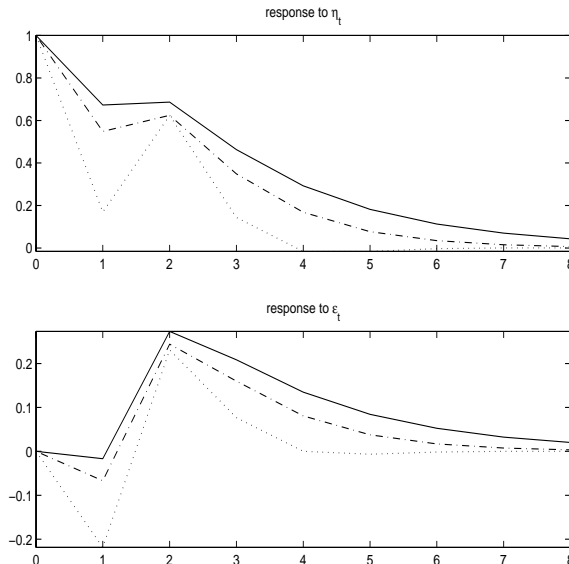


FIGURE 9. Top panel: impulse responses of inflation to the shock η_t under three values of θ : $\theta = +\infty$ (solid line), $\theta = 10$ (dashed-dotted line), and $\theta = 5$ (dotted line), with $\beta = 1$. Bottom panel: impulse response of inflation to shock ϵ_t under same three values of θ .

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 8 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 criterion (23). Flattening of 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 feedback rule $r_t = ay_t + b\pi_t$.²³ This effect of activating a preference for robust rules has the following interpretation. Ball's model specifies that the shocks in (10),(11), (12) are serially uncorrelated. Under the optimal ($\theta = +\infty$) rule, this rule would expose the policy maker to the biggest costs were the shocks instead to be highly positively serially correlated. Under a preference for robustness, the policy maker is most concerned about misreading a 'permanent' or 'long-lived' shock as a temporary (i.e., serially uncorrelated) one. To protect himself, he responds to serially uncorrelated shocks (under the reference model) as though they were positively serially correlated. This manifests itself in making the interest rate more responsive to both y_t and π_t .

²³See (Sargent 1999a) for a discussion.

An interesting aspect of the two panels of Figure 8 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 9 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 figures show that activating preferences for robustness causes the impulse responses to damp out more quickly, which is consistent with the flatter $[\text{trace } h(\omega)'h(-\omega)]$'s 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 complaint about 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 mechanism 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.

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. When economists study risk aversion, they have in mind confronting decision makers with gambles with known probabilities. From knowledge or guesses about how people would behave when confronted with specific and well defined risks, economists infer risk aversion that is *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 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). Historical evidence from security market data suggest that risk aversion implied by security market data is very much larger than that elicited from participants in hypothetical gambles.

There is a variety of responses to this puzzle. They range from questioning the empirical measurements of the risk-return trade-off to challenging the ability to extrapolate risk aversion from different-sized gambles. Indeed, mean returns on equity are known to be difficult to measure reliably. Also, it has been claimed that people look more risk averse when facing small gambles than large ones. ((Segal and Spivak 1990), (Epstein and Melino 1995) and (Rabin 1999)). Our statistical notion of robustness makes contact with the first of these explanations.

A preference for 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 (1995) build on work of Jacobson (1973) and Whittle (1980) to establish an equivalence between a preference for robustness and risk-sensitive preferences for game 3. Anderson, Hansen, and Sargent (1999) extend this equivalence result to a larger class of recursive robustness 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.²⁴

Hansen, Sargent, and Tallarini (1999), Tallarini (2000), Cagetti, Hansen, Sargent, and Williams (2000) 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 to detect in the behavior of 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 equilibrium market prices of risk. Hansen, Sargent, and Tallarini (1999) and Hansen, Sargent, and Wang (2000) 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 or reference 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

²⁴This 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 the game 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 10. 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.

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) and Chen and Epstein (1998). One 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 (1999) 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 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 10 is from Hansen, Sargent, and Wang (2000), which studied the contribution to the market price of risk from a preference for 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 10 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. For a given detection error probability, the value of θ depends on the model. (See the discussion above about how the the detection error probability depends on θ and the particular model.) Figure 10 affirms the tight link between detection error probabilities and the contribution of a preference for robustness to the market price of risk that was asserted by Anderson, Hansen, and Sargent (1999). 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 game in which a malevolent nature chooses the model as an instrument to achieve robustness across the set of models.

In the body of the paper, we did not touch 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, Sargent, and Tallarini (1999) and Hansen and Sargent (2000) use an equilibrium concept that seems a natural extension of rational expectations where all agents share the same set of alternative models (i.e., the same θ in the notation above). Suitably viewed, the communism of models seen in rational expectations models extends to this setting: now agents share an approximating model, a set of surrounding models against which they value robustness, and also a synthesized worst case model. However, even in settings in which agents share an approximating model, if their tastes for robustness differ – that is, if their sets of surrounding models differ – then the equilibrium must be modified. Anderson (1998) studies a pure endowment economy whose agents have what we would interpret as different preferences for robustness, and shows how the distribution of wealth over time is affected by those preferences.²⁵ Hansen and Sargent (2000) and Kasa (1999) describe multi-agent problems in the form of Ramsey problems for a government facing a competitive private economy.

²⁵ Anderson 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.

Preferences for robustness also bear on the Lucas critique. The Lucas 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 each shock process, a rational expectations theory associates a distinct decision rule. Lucas criticised 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, apparently ‘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.

REFERENCES

- Anderson, E. (1998). Uncertainty and the dynamics of pareto optimal allocations. University of Chicago Dissertation.
- 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. *Handbook in Computational Economics*, H. M. Amman, D. A. Kendrick and J. Rust, editors, 171–252.
- Ball, L. (1999). Policy rules for open economies. In J. Taylor (Ed.), *Monetary Policy Rules*, pp. 127–144. University of Chicago Press.
- 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.
- Brainard, W. (1967). Uncertainty and the effectiveness of policy. *American Economic Review* 57, 411–425.
- 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 (2000). Robustness and filtering in a stochastic growth model. Mimeo.
- Camerer, C. (1995). Individual decision making. *The Handbook of Experimental Economics*, edited by John H. Kagel and Alvin Roth, 588–673.
- 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*, 115–169.
- 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.
- 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., 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 (1995). Discounted linear exponential quadratic gaussian control. *IEEE Transactions on Automatic Control* 40, 968–971.
- Hansen, L. P. and T. J. Sargent (2000, February). Elements of robust filtering and control for macroeconomics. Mimeo.
- Hansen, L. P., T. J. Sargent, and N. E. Wang (2000). Robust permanent income and pricing with filtering. Mimeo.
- 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.
- 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 E. K. Donald P. Jacobs 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.
- 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 (1997). Recurrent hyperinflations and learning. Universitat Pompeu Fabra.
- Miller, B. L. (1974). Optimal consumption with a stochastic income stream. *Econometrica* 42, 253–266.
- Onatski, A. and J. H. Stock (1999). Robust monetary policy under model uncertainty in a small model of the u.s. economy. Mimeo.
- Rabin, M. (1999, January). Risk aversion and expected-utility theory: A calibration theorem. Mimeo.
- 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.
- 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.
- 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, STANFORD UNIVERSITY, STANFORD, CALIFORNIA

E-mail address: sargent@leland.stanford.edu