Review of Economic Dynamics **4,** 519–535 (2001)
doi:10.1006/redy.2001.0132, available online at http://www.idealibrary.com on IDEAL®

Acknowledging Misspecification in Macroeconomic Theory[1]

Lars Peter Hansen

Department of Economics, University of Chicago, Chicago, Illinois
E-mail: l-hansen@uchicago.edu

and

Thomas J. Sargent

Department of Economics and Hoover Institution, Stanford University, Stanford, California
E-mail: sargent@leland.stanford.edu

Received January 25, 2001

We explore methods for confronting model misspecification in macroeconomics. We construct dynamic equilibria in which private agents and policy makers recognize that models are approximations. We explore two generalizations of rational expectations equilibria. In one of these equilibria, decision-makers use dynamic evolution equations that are imperfect statistical approximations, and in the other misspecification is impossible to detect even from infinite samples of time series data. In the first of these equilibria, decision rules are tailored to be robust to the allowable statistical discrepancies. Using frequency domain methods, we show that robust decision-makers treat model misspecification like time series econometricians.

Key Words: robustness; model misspecification; monetary policy; rational expectations.

[1] This paper was prepared for the Society of Economic Dynamics Conference in Costa Rica in June 2000, for the Ninth International Conference of the Institute for Monetary and Economic Studies, Bank of Japan Conference in Tokyo on July 3–4 in 2000, and for Sargent's Frisch lecture at the World Congress of the Econometric Society at Seattle in August 2000. We thank Gary Hansen, Lars Svensson, Grace Tsiang, and two referees for this journal for helpful comments.

1094-2025/01 $35.00

1. RATIONAL EXPECTATIONS VERSUS MISSPECIFICATION

Subgame perfect and rational expectations equilibrium models do not permit a self-contained analysis of model misspecification. But sometimes model builders suspect misspecification, and so might the agents in their model.[2] To study this we must modify rational expectations. But in doing so, we want to respect and extend the inspiration underlying rational expectations, which was to deny that a model builder knows more about the data generating mechanism than do the agents inside his model.

This paper describes possible reactions of model builders and agents to two different types of model misspecification. The first type is difficult to detect in time series samples of the moderate sizes typically at our disposable. A second type of model misspecification is impossible to detect even in infinite samples drawn from an equilibrium.

1.1. *Rational Expectations Models*

A model is a probability distribution over a sequence. A rational expectations equilibrium is a fixed point of a mapping from agents' personal models of an economy to the actual model. Much of the empirical power of rational expectations models comes from identifying agents' models with the data generating mechanism. Leading examples are the cross-equation restrictions coming from agents' using conditional expectations to forecast and the moment conditions emanating from Euler equations. A persuasive argument for imposing rational expectations is that agents have incentives to revise their personal models to remove readily detectable gaps between them and the empirical distributions. The rational expectations equilibrium concept is often defended as the limit point of some more or less explicitly specified learning process in which all personal probabilities eventually merge with a model's population probability.

1.2. *Recognitions of Misspecification*

A rational expectations equilibrium (indexed by a vector of parameters) *is* a likelihood function. Many authors of rational expectations models express or reveal concerns about model misspecification by declining to use the model (i.e., the likelihood function) for empirical work. One example is the widespread practice of using seasonally adjusted and/or low-frequency adjusted data. Those adjustments have been justified formally by stressing the model's inadequacy at particular frequencies and by

[2] By studying how agents who fear misspecification can promote cautious behavior and boost market prices of risk, we do not intend to deny that economists have made tremendous progress by using equilibrium concepts that ignore model misspecification.

appealing to some frequency domain version of an approximation criterion like that of Sims (1972), which is minimized by least squares estimates of a misspecified model. Sims (1993) and Hansen and Sargent (1993) describe explicit justifications that distinguish the model from the unknown true data generating process. They posit that the true generating process has behavior at the seasonal frequencies that cannot be explained by the model except at parameter values that cause bad fits at the nonseasonal frequencies. Maximum likelihood estimates make the model best fit the frequencies contributing the most variance to the data set. When the model is most poorly specified at the seasonal frequencies, using seasonally adjusted data can trick the maximum likelihood method to emphasize frequencies where the model is better specified. This can give better estimates of parameters describing tastes and technologies.[3]

Less formal reasons for divorcing the data analysis from the model also appeal to model misspecification. For example, calibrators say that their models are approximations aimed at explaining only "stylized" facts or particular features of the time series.

Notice that in both the formal defenses of data filtering and the informal practice of calibration, the economist's model typically remains a rational expectations model inhabited by agents who do not doubt the model. Thus, such analyses do not let the agents inside the economist's model share his doubts about model specification.

2. AGENTS WHO SHARE ECONOMISTS' DOUBTS

But the intent of rational expectations is to put the economist and the agents inside his model on the same footing. Letting the agents contemplate model misspecification reopens fundamental issues that divided Knight, Fellner, and Ellsberg from Savage and that were set aside when, by adopting rational expectations, macroeconomists erased all model ambiguity from their agents' minds.

2.1. *Savage versus Knight*

Knight (1921) distinguished risky events, which could be described by a probability distribution, from a worse type of ignorance that he called uncertainty that could not be described by a probability distribution. He

[3] The remarkable feature of these results is that better estimates of taste and technology parameters are acquired by imposing *false* cross-equation restrictions and by accepting *worse* estimates of the parameters governing information and agents' forecasts. Two sided seasonal adjustment distorts the temporal and information properties processes that agents are trying to forecast.

thought that profits compensated entrepreneurs for bearing uncertainty. Especially in some urn examples that prefigured Ellsberg (1961), we see Knight thinking about decision making in the face of possible model misspecifications.[4] Savage contradicted Knight. Savage (1954) proposed a set of axioms about behavior that undermined Knight's distinction between risk and uncertainty. A person behaving according to Savage's axioms has a well-defined personal probability distribution that rationalizes his behavior as an expected utility maximizer. Savage's system undermined Knight by removing the agent's possible model misspecification as a concern of the model builder.

2.2. *Muth versus Savage*

For Savage, it was not an aspect of rationality that personal probabilities be "correct." But for followers of Muth (1961), it was. By equating personal probabilities with objective ones, rational expectations assumes away possible model misspecifications and disposes of diversity of personal probabilities. Rational expectations substantially weaken the appeal of Savage's "solution" of the model specification problems that concerned Knight because they so severely restrict personal probabilities.

2.3. *Ellsberg versus Savage*

On the basis of experimental evidence, Ellsberg (1961) and Fellner (1961) challenged Savage's theory. Fellner (1965) proposed a semiprobabilistic framework in which agents use context-specific "slanted probabilities" to make decisions in ways that violate the Savage axioms. The Ellsberg paradox motivated Gilboa and Schmeidler (1989) to formulate a new set of axioms that accommodate model ambiguity. Gilboa and Schmeidler's axioms give agents not a unique personal probability distribution but a *set* of distributions. They posit that agents make decisions as the max–min outcomes of a two-person game in which the agent chooses a utility maximizing decision and a malevolent nature chooses a minimizing probability distribution from within the set chosen by the agents. They show that such behavior can explain the Ellsberg paradox.

Convinced by the Ellsberg paradox and inspired by Gilboa and Schmeidler's formulation, Epstein and Wang (1994), Epstein and Melino (1995), and Chen and Epstein (1998) have constructed dynamic models in which agents are adverse to model ambiguity. Some of this work represents model ambiguity by a class of probability distributions generated by the epsilon-contaminations used in the robust statistics literature.

[4] As Ellsberg (1961) pointed out, Knight's introspection about urns did not produce the paradox that Ellsberg is famous for.

2.4. *Brunner and Meltzer*

Brunner and Meltzer (1967) discussed the role of model misspecification in the design of monetary policy. They challenged the existence of a "fully identified, highly confirmed theory of macroeconomic processes." They wrote:

> ...we acknowledge that many of the questions raised here can be answered more fully if (and only if) more useful knowledge about the structure of the economy is assumed or obtained. Put otherwise, the theorist may choose to ignore this problem by assuming the possession of reliable information currently outside the scope of quantitative economics. The policy maker is not as fortunate.[5]

By way of acknowledging policy makers' model ambiguity, Brunner and Meltzer advocated a min−max strategy for selecting among endogenous indicators of monetary policy.

We share Brunner and Meltzer's concern about model uncertainty. Starting from a single dynamic model, we add perturbations that represent potential model misspecifications around that benchmark model. The perturbations can be viewed as indexing a large family of dynamic models, as in dynamic extensions of the Gilboa-Schmeidler multiple prior formulation. We prefer to think about the perturbations as errors in a convenient, but misspecified, dynamic macroeconomic model. We take the formal structure of our model perturbations from a source that served macroeconomists well before, especially at the dawn of rational expectations.

3. CONTROL THEORY

The mathematical literatures on control and estimation theory were the main sources of tools for building rational expectations models. That was natural, because before 1975 or so, control theory was about how to design optimal policies under the assumption that a decision maker's model, typically a controlled Markov process, is correctly specified. This is exactly the problem that rational expectations modelers want the agents inside their models to solve.

But just when macroeconomists were importing the ordinary correct-model version of control theory to refine rational expectations models, control theorists were finding that policies they had designed under the assumption that a model is correct sometimes performed much worse than they should. Such practical problems caused control theorists to soften their assumption about knowing the correct model and to think about ways

[5] See Brunner and Meltzer (1967), page 188.

to design policies for models that were more or less good approximations. Starting in the mid-1970's, they created tools for designing policies which would work well for a set of possible models that were in a sense close to a basic approximating model. In the process, they devised manageable ways of formulating a set of models surrounding an approximating model. To design policies in light of that set of models, they used a min–max approach like that of Gilboa and Schmeidler.

4. ROBUST CONTROL THEORY

We want to import, adapt, and extend some of the robust control methods to build models of economic agents who experience model ambiguity. We briefly sketch major components of our analysis.

4.1. *The Approximating Model*

There is a single approximating model defining transition probabilities. For example, consider the linear–quadratic state space system

$$x_{t+1} = Ax_t + Bu_t + Cw_{t+1} \tag{1}$$

$$z_t = Hx_t + Ju_t, \tag{2}$$

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

$$-\frac{1}{2}\sum_{t=0}^{\infty}\beta^t E|z_t|^2, \tag{3}$$

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 $u_t = -Fx_t$.

4.2. *Distortions*

A particular way of distorting those probabilities defines a set of models that express an agent's model ambiguity. First, we need a concise way to describe a class of alternative specifications. For a given policy $u = -Fx$, the state equation (1) defines a Markov transition kernel $\pi(x' \mid x)$. For all x, let $v(x' \mid x)$ be positive for all x. We can use $v(x' \mid x)$ to define a distorted model via $\pi^v(x' \mid x) = \frac{v(x' \mid x)\pi(x' \mid x)}{Ev(x' \mid x)}$, where the division by $Ev(x' \mid x)$ lets the quotient be a probability. Keeping $v(x' \mid x)$ strictly positive means that the two models are mutually absolutely continuous, which makes the models difficult to distinguish in small data sets. *Conditional relative entropy* is a measure of the discrepancy between the approximating model and the distorted model. It is defined as

$$\text{ent} = \int \log \frac{\pi^{v}(x' \mid x)}{\pi(x' \mid x)} \pi^{v}(dx' \mid x). \tag{4}$$

Conditional entropy is thus the conditional expectation of the log likelihood ratio of the distorted to the approximating model, evaluated with respect to the distorting model.

The distortion just described preserves the first-order Markov nature of the model. This occurs because in $v(x' \mid x)$ only x shows up in the conditioning. More general distortions allow lags of x to show up in the conditioning information set. To take a specific example, we represent perturbations to model (1) by distorting the conditional mean of the shock process away from zero:

$$x_{t+1} = Ax_t + Bu_t + C(w_{t+1} + v_t) \tag{5}$$

$$v_t = f_t(x_t, \dots, x_{t-n}). \tag{6}$$

Here Eq. (6) is a set of distortions to the conditional means of the innovations of the state equation. In (6), these are permitted to feed back on lagged values of the state and thereby represent misspecified dynamics. For the particular model of the discrepancy (6), it can be established that

$$\text{ent} = \frac{1}{2} v_t' v_t.$$

This distortion substitutes a higher order nonlinear Markov model for the first-order linear approximating model. More generally, the finite-order Markov structure can be relaxed by supposing that v_t depends on the infinite past of x_t. The essential requirement is the mutual absolute continuity between the approximating model and its perturbed counterpart.

4.3. *Conservative Valuation*

We use a min–max operation to define a conservative way of evaluating continuation utility. Let $V(x_{t+1})$ be a continuation value function. Fix a control law $u_t = -Fx_t$ so that the transition law under a distorted model becomes

$$x_{t+1} = A_o x_t + C(w_{t+1} + v_t), \tag{7}$$

where $A_o = A - BF$. We define a distorted expectations operator $R_\theta(V(x_{t+1}))$ for evaluating continuation values as the indirect utility function for the minimization problem

$$R_\theta(V(x_{t+1})) = \min_{v_t} \{\theta v_t' v_t + E(V(x_{t+1}))\}, \tag{8}$$

where the minimization is subject to constraint (7). Here $\theta \leq +\infty$ is a parameter that penalizes the entropy between the distorted and approximating models and thereby describes the size of the set of alternative models for which the decision maker wants a robust rule. This parameter is context-specific and depends on the confidence in his model that the historical data used to build it inspire. Below, we illustrate how detection probabilities for discriminating between models can be used to discipline the choice of θ.

Let $\hat{v}_t = G_\theta x_t$ attain the minimum on the right side. The following useful robustness bound follows from the minimization on the right side of Eq. (8):

$$EV(A_o + C(w_{t+1} + v_t)) \geq R_\theta(V(x_{t+1})) - \theta v_t' v_t. \tag{9}$$

The left side is the expected continuation value under a distorted model. The right side is a lower bound on that continuation value. The first term is a conservative value of the continuation value under the approximating ($v_t = 0$) model. The second term on the right gives a bound on the rate at which performance deteriorates with the entropy of the misspecification. Decreasing the penalty parameter θ lowers the $R_\theta V(x_{t+1})$, thereby increasing the conservative nature of $R_\theta V(x_{t+1})$ as an estimate under the approximating model and also decreasing the rate at which the bound on performance deteriorates with entropy. Thus, the indirect utility function induces a conservative way of evaluating continuation utility that can be regarded as a context-specific distortion of the usual conditional expectation operator. There is a class of operators indexed by a single parameter that summarizes the size of the set of alternative models.

It can be shown that the operator R_θ can be represented as

$$\begin{aligned} R_\theta(V(x_{t+1})) &\approx h^{-1} E_t(h(V(x_{t+1})) \\ &\approx \theta \hat{v}_t' \hat{v}_t + E_t(V(x_{t+1})), \end{aligned}$$

where $h(V) = -\exp(\frac{-V}{\theta})$ and $\hat{v}_t$ is the minimizing choice of v_t.[6] Note that h is a concave function. This is an operator used by Epstein and Zin (1989), Weil (1993), and others, but with no connection to robustness.

The operator R_θ can be used to rework the theory of asset pricing or to design robust decision rules. The operator also has been used by Hansen and Sargent (1995) to define a discounted version of risk-sensitive preference specification of Jacobson (1973) and Whittle (1990). They parameterize risk-sensitivity by a parameter $\sigma \equiv -\theta^{-1}$, where the $\sigma < 0$ imposes an additional adjustment for risk that acts like a preference for robustness.

4.4. *Robust Decision Rule*

A robust decision rule is produced by the Markov-perfect equilibrium of a two-person zero-sum game in which a maximizing agent chooses a policy and a minimizing agent chooses a model. To compute a robust control rule we use the Markov-perfect equilibrium of the two-agent zero-sum dynamic game,[7]

$$V(x) = \max_u \min_v \left\{ -\frac{1}{2} z'z + \frac{\beta\theta}{2} v'v + \beta E_t V(x^*) \right\} \tag{10}$$

subject to

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

Here $\theta > 0$ is a penalty parameter that constrains the minimizing agent; θ governs the degree of robustness achieved by the associated decision rule. When the robustness parameter θ takes the value $+\infty$, we have ordinary control theory because it is too costly for the minimizing agent to set a nonzero distortion v_t. Lower values of θ achieve some robustness (again see the role of θ in the robustness bound (9)).

To illustrate robustness we present some figures from Hansen and Sargent (2000b) based on a monetary policy model of Ball (1999).[8] Ball's is a "backward looking" macro model with the structure

$$y_t = -\beta r_{t-1} - \delta e_{t-1} + \epsilon_t \tag{11}$$

$$\pi_t = \pi_{t-1} + \alpha y_{t-1} - \gamma(e_{t-1} - e_{t-2}) + \eta_t \tag{12}$$

$$e_t = \theta r_t + \nu_t, \tag{13}$$

[6] We use the notation $\approx$ because there is a difference in the constant term that becomes small when we take a continuous-time diffusion limit.

[7] There is only one value function because it is a zero-sum game.

[8] See Hansen and Sargent (2001) for a parallel but more extended numerical exercise that studies robust decision rules for a single-agent permanent income model. Robustness gives rise to a form of precautionary savings that depends on the magnitude of variances of innovations to income but that does not require convexity of marginal continuation values.

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, and ν are serially uncorrelated and mutually orthogonal disturbances. As an objective, Ball assumed that the monetary authority wants to maximize

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

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

Ball motivated (11) as an open-economy IS curve and (12) as an open-economy Phillips curve; he uses (13) to capture effects of the interest rate on the exchange rate. Ball set the parameters γ, θ, β, δ at the values 0.2, 2, 0.6, 0.2. Following Ball, we set the innovation shock standard deviations equal to 1, 1, $\sqrt{2}$.

Ball's model is particularly simple because the private sector is "backward-looking." This reduces the monetary policy problem to a single agent decision problem, a simple control problem. While this simplifies our representation of robustness, it would be of more substantive interest to investigate models in which private sector decision-makers are "forward-looking." Hansen and Sargent (2001) adapt ideas of Currie and Levine (1987), Pearlman (1992), and Pearlman *et al.* (1986) to study models in which both the Federal Reserve and private agents are forward-looking and in which both are concerned about model misspecification. It is possible to compute Ramsey plans for models with forward looking agents by solving a single Bellman equation of the form (10). Thus, from a technical viewpoint, it is tractable to design policy rules for forward looking models using our tools.

Under various worst-case models indexed by the value of $\sigma = \theta^{-1}$ on the horizontal axis, Fig. 1 shows the values of $-Ez_t'z_t$ for the three rules that we have labeled with values of $\sigma = 0, -0.04, -0.085$. Later we shall describe how these different settings of $\sigma = -\theta^{-1}$ correspond to different sizes of the set of alternative models for which the decision maker seeks robustness. For Ball, $-Ez_t'z_t$ is set equal to $-E(\pi_t^2 + y_t^2)$, which equals minus the sum of variances of inflation and output. The three fixed rules solve Eq. (10) for the indicated value of σ. The value of $-Ez_t'z_t$ is plotted for each of the fixed rules evaluated for the law of motion $x_{t+1} = (A - BF)x_t + C(w_{t+1} + G(\tilde{\sigma})x_t)$ where $G(\tilde{\sigma})x_t$ denotes the minimizing rule for v_t associated with the value $\sigma = \tilde{\sigma}$ on the horizontal axis. The way the curves cross indicates how the $\sigma = -0.04$ and $\sigma = -0.085$ rules are not optimal if the model is specified correctly (if $\sigma = 0$ on the horizontal axis), but do better than the optimal rule against the model misspecifications associated with the distortions associated with movements along the horizontal axis. Notice how the σ used to design the rule affects the slope of the payoff line. We now briefly turn to describe how σ might be chosen.

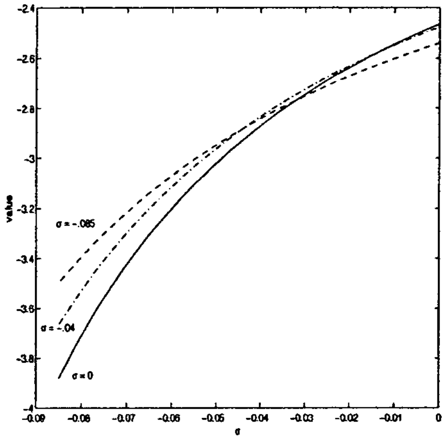

FIG. 1. Value of $-Ez_t'z_t$ for three decision rules when the data are generated by the worst-case model associated with the value of σ on the horizontal axis: $\sigma = 0$ rule (solid line), $\sigma = -.04$ rule (dashed–dotted line), $\sigma = -.085$ (dashed) line. The robustness parameter is given by $\theta = -1/\sigma$.

4.5. *Detection Probabilities*

The Bellman equation (10) specifies and penalizes model distortions in terms of the conditional relative entropy of a distorted model with respect to an approximating model. Conditional relative entropy is a measure of the discrepancy between two models that appears in the statistical theory of discriminating between two models. We can link detection error probabilities to conditional relative entropy, as in Anderson *et al.* (1999) and Hansen *et al.* (2000). This allows us to discipline our choice of the single free parameter θ that our approach brings relative to rational expectations.

For a sample of 147 observations, Fig. 2 displays a set of Bayesian detection error probabilities for comparing Ball's model with the worst case model from Eq. (10) that is associated with the value of $\sigma = -\theta^{-1}$ on the axis. The detection error probability is .5 for $\sigma = 0$ (Ball's model and the $\sigma = 0$ worst-case model are identical and therefore detection errors occur half the time). As we lower σ, the worst-case model from (10) diverges more and more from Ball's model (because $v_t'v_t$ rises) and the detection error probability falls. For $\sigma = -0.04$, the detection error probability is still 0.25: this high fraction of wrong judgments from a model comparison test tells us that it is difficult to distinguish Ball's model from the worst-case $\sigma = -0.04$ model with 147 observations. Therefore, we think it is reasonable for the monetary authority in Ball's model to want to be robust against misspecifications parameterized by such values of σ. In

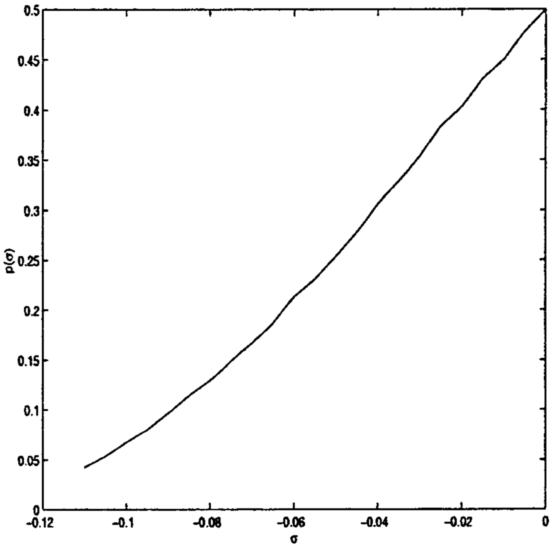

FIG. 2. Detection error probability (coordinate axis) as a function of $\sigma = -1/\theta$ for Ball's model.

this way, we propose to use a table of detection error probabilities like that encoded in Fig. 2 to discipline our selection of σ. See Anderson *et al.* (1999) and Hansen *et al.* (2000) for applications to asset pricing.

4.6. *Precaution*

A preference for robustness induces context-specific precaution. In asset pricing models, this boosts market prices of risk and pushes a model's predictions in the direction of the data with respect to the equity premium puzzle. See Hansen *et al.* (1999) and Hansen *et al.* (2000). In permanent income models, it induces precautionary savings. In sticky-price models of monetary policy, it can induce a policy authority to be more aggressive in response to shocks than one who knows the model. Such precaution has an interpretation in terms of a frequency domain representation of the criterion function under various model perturbations.

For the complex scalar ζ, let $G(\zeta)$ be the transfer function from the shocks w_t to the targets z_t. Let $'$ denote matrix transposition and complex conjugation, $\Gamma = \{\zeta : |\zeta| = \sqrt{\beta}\}$, and $d\lambda(\zeta) = (1/2\pi i\sqrt{\beta}\,\zeta)d\zeta$. Then the criterion (3) can be represented as

$$H_2 = -\int_\Gamma \operatorname{trace}[G(\zeta)'G(\zeta)]\,d\lambda(\zeta).$$

Where $-\delta_j(\zeta)$ is the jth eigenvalue of $G(\zeta)'G(\zeta)$, we have

$$H_2 = \sum_j \int_\Gamma -\delta_j(\zeta)\, d\lambda(\zeta). \tag{14}$$

Hansen and Sargent (2001) show that a robust rule is induced by using the criterion

$$\mathrm{ent}(\theta) = \int_\Gamma \log\det[\theta I - G(\zeta)'G(\zeta)]\, d\lambda(\zeta)$$

or

$$\mathrm{ent}(\theta) = \sum_j \int_\Gamma \log\big[\theta - \delta_j(\lambda)\big]\, d\lambda(\zeta). \tag{15}$$

Because $\log(\theta - \delta)$ is a concave function of $-\delta$, criterion (15) is obtained from criterion (14) by putting a concave transformation inside the integration. Aversion to model misspecification is thus represented as additional "risk aversion" across frequencies instead of across states of nature. Under criterion (15), the decision maker prefers decision rules that render trace $G(\zeta)'G(\zeta)$ flat across frequencies.

For example, the curve for $\sigma = 0$ in Fig. 3 depicts a frequency domain decomposition of $Ez_t'z_t$ for the optimal rule under Ball's model. Notice that it is greatest at low frequencies. This prompts the minimizing agent in problem (10) to make what Ball's model assume to be i.i.d. shocks instead be serially correlated (by making v_t feed back appropriately on x_t). The maximizing agent in (10) responds by changing the decision rule so that it is less vulnerable to low-frequency misspecifications. In Ball's model, this can be done by having the monetary authority adjust interest rates more aggressively in its "Taylor rule." Notice how the frequency decompositions under the $\sigma = -.04$ and $\sigma = -.085$ rules are flatter.[9] They thereby achieve robustness by rendering themselves less vulnerable to low-frequency misspecifications.

A permanent income model is also most vulnerable to misspecifications of the income process at the lowest frequencies, since it is designed to do a good job at smoothing high frequency movements. For a permanent income model, a preference for robustness with respect to the specification of the income process then induces precautionary saving of a type that does not depend on the third derivative of the value function.

[9] These are frequency decompositions of the Ball's criterion function operating under the robust rules when the approximating model governs the data.

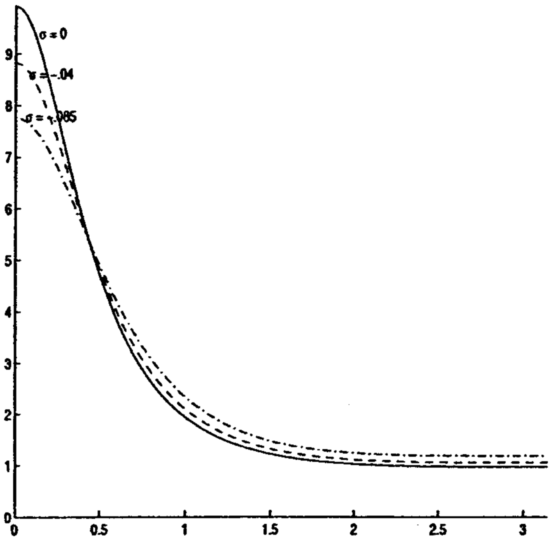

FIG. 3. Frequency decompositions of $Ez_t'z_t$ for objective function of Ball's model under three decision rules ($\sigma = 0, -0.04, -0.085$). The robustness parameter satisfies $\theta = -1/\sigma$.

4.7. *Multiagent Settings*

We have discussed only single-agent decision theory, despite the presence of the minimizing second agent. The minimizing agent in the Bellman equation (10) is fictitious, a computational device for the maximizing agent to attain a robust decision rule. Macroeconomists routinely use the idea of a representative agent to study aggregate phenomena using single-agent decision theory, for example by studying planning problems and their decentralizations. We can use a version of Eq. (10) in conjunction with such a representative agent device. A representative agent formulation would attribute a common approximating model and a common set of admissible model perturbations to the representative agent and the government. Robust Ramsey problems can be based on versions of problem (10) augmented with implementability constraints. See Hansen and Sargent (2001) and Hansen and Sargent (2000a) for some examples.

5. SELF-CONFIRMING EQUILIBRIA

We have focused on misspecifications that are difficult to detect with moderate-sized data sets, but that can be distinguished with infinite ones. We now turn to a more subtle kind of misspecification, one that is beyond the capacity of detection error probabilities to unearth even in infinite samples. It underlies the concept of self-confirming equilibrium, a type of rational expectations that seems natural for macroeconomics. A self-confirming equilibrium attributes possibly distinct personal probabilities (models) to each agent in the model. Those personal probabilities (1) are permitted to differ on events that occur with zero probability at equilibrium, but (2) must agree on events that occur with positive probability at equilibrium. Requirement (2) means that the differences among personal probabilities cannot be detected even from infinite samples from the equilibrium. Fudenberg and Kreps (1995a), Fudenberg and Kreps (1995b), Fudenberg and Levine (1993), and Sargent (1999) advocate the concept of self-confirming equilibrium partly because it is the natural limit point of a set of adaptive learning schemes (see Fudenberg and Levine (1998)). The argument that agents will eventually adapt to eliminate discrepancies between their model and empirical probabilities strengthens the appeal of a self-confirming equilibrium but does nothing to promote subgame perfection.[10] A self-confirming equilibrium is a type of rational expectations equilibrium. However, feature (1) permits agents to have misspecified models that fit the equilibrium data as well as any other model but that miss the "causal structure." The beliefs of a large player about what will occur off the equilibrium path can influence his choices and therefore outcomes along the equilibrium path.

That self-confirming equilibria permit large players—in particular governments—to have wrong views provides ways to resolve what we at Minnesota used to call "Wallace's conundrum" in the mid-1970's.[11] Wallace had in mind a subgame perfect equilibrium. He noted that there is no room for policy advice in a rational expectations equilibrium where private agents know the conditional probabilities of future choices of the government. In such a (subgame perfect) equilibrium, there are no free variables for government agents to choose: their behavior rules are already known and responded to by private agents. For example, if a researcher believes that the historical data obey a Ramsey equilibrium for some dynamic optimal tax problem, he has no advice to give other than maybe to change the environment.

[10] Bray and Kreps (1987) present a Bayesian analysis of equilibrium learning without misspecification. They express regret that they precluded an analysis of learning *about* a model by having set things up so that their agents can be Bayesians (which means that they know the model from the beginning and learn only *within* the model). Agents' model misspecification, which disappears only eventually, is a key part of how Fudenberg and Levine (1998) and others analyze learning about an equilibrium.

[11] See Sargent and Wallace (1976).

Self-confirming equilibria contain some room for advice based on suggestions for improving the government's model.[12] But such advice is likely to be resisted because the government's model fits the historical data as well as the critic's model. Therefore, those who criticize the government's model must do so on purely theoretical grounds, or else wait for unprecedented events to expose an inferior fit of the government's model.

REFERENCES

Anderson, E., Hansen, L. P., and Sargent, T. (1999). Robustness, detection and the price of risk, mimeo.

Ball, L. (1999). Policy rules for open economies, in *Monetary Policy Rules* (J. Taylor, Ed.), Chicago: Univ. of Chicago Press.

Bray, M. M., and Kreps, D. M. (1987). Rational learning and rational expectations, in *Arrow and the Ascent of Modern Economic Theory*, (G. R. Feiwel, Ed.), New York: New York Univ. Press.

Brunner, K. and A. H. Meltzer (1967). The meaning of monetary indicators, in *Monetary Process and Policy: A Symposium* (G. Horwich, Ed.), Homewood, IL: Irwin.

Chen, Z. and L. Epstein (1998). Ambiguity, risk, and asset returns in continuous time, mimeo.

Currie, D. and P. Levine (1987). The design of feedback rules in linear stochastic rational expectations models, *Journal of Economic Dynamics and Control* **11**(1), 1–28.

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.

Epstein, L. G. and S. Zin (1989). Substitution, risk aversion, and the temporal behavior of asset returns: A theoretical framework, *Econometrica* **57**, 937–969.

Fellner, W. (1961). Distortion of subjective probabilities as a reaction to uncertainty, *Quarterly Journal of Economics* **75**, 670–689.

Fellner, W. (1965). *Probability and Profit: A Study of Economic Behavior along Bayesian Lines*, Homewood, IL: Irwin.

Fudenberg, D. and D. M. Kreps (1995a). Learning in extensive games. I. Self-confirming and Nash equilibrium, *Games and Economic Behavior* **8**, 20–55.

Fudenberg, D. and D. M. Kreps (1995b). Learning in extensive games. II. Experimentation and Nash equilibrium, mimeo, Harvard University.

Fudenberg, D. and D. K. Levine (1993). Self-confirming equilibrium. *Econometrica* **61**, 523–545.

Fudenberg, D. and D. K. Levine (1998). *The Theory of Learning in Games*, Cambridge, MA: MIT Press.

[12] Much of the macroeconomists' advice has been of that form: think of the arguments about the natural unemployment rate theory, which were about using new cross-equation restrictions to reinterpret existing econometric relations.

Gilboa, I., and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior, *Journal of Mathematical Economics* **18**, 141–153.

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 (2000a). Robust control and filtering in forward-looking models, unpublished manuscript.

Hansen, L. P. and T. J. Sargent (2000b). Wanting robustness in macroeconomics, unpublished manuscript.

Hansen, L. P. and T. J. Sargent (2001). Robust control and filtering 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.

Knight, F. H. (1921). *Risk*, *Uncertainty and Profit*, Boston/New York: Houghton Mifflin.

Muth, J. F. (1961). Rational expectations and the theory of price movements, *Econometrica* **29**(3), 315–335.

Pearlman, J. G. (1992). Reputational and nonreputational policies under partial information, *Journal of Economic Dynamics and Control* **16**(2), 339–358.

Pearlman, J. G., D. A. Currie, and P. L. Levine (1986). Rational expectations models with partial information, *Economic Modeling* **3**(2), 90–105.

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

Sargent, T. J. and N. Wallace (1976). Rational expectations and the theory of economic policy. *Journal of Monetary Economics* **2**(2), 169–183.

Savage, L. J. (1954). *The Foundations of Statistics*, New York: Wiley.

Sims, C. A. (1972). Approximate prior restrictions in distributed lag estimation, *Journal of the American Statistical Association* **67**, 169–175.

Sims, C. A. (1993). Rational expectations modeling with seasonally adjusted data, *Journal of Econometrics* **55**, 9–19.

Weil, P. (1993). Precautionary savings and the permanent income hypothesis, *Review of Economic Studies* **60**, 367–383.

Whittle, P. (1990). *Risk-Sensitive Optimal Control*, New York: Wiley.