Journal of Economic Dynamics & Control 34 (2010) 1951–1966

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

Robust hidden Markov LQG problems

Lars Peter Hansen[a], Ricardo Mayer[a,b], Thomas Sargent[c,d,*]

[a] The University of Chicago, USA
[b] Universidad Diego Portales, Chile
[c] New York University, 19 W. 4th Street, 6FL, New York, NY 10012, USA
[d] Hoover Institution, Stanford University, USA

ARTICLE INFO

Article history:
Received 15 July 2009
Accepted 7 April 2010
Available online 11 May 2010

JEL classification:
C11
C12
C61
E61

Keywords:
Hidden Markov models
Misspecification
Kalman filter
Robustness
Entropy
Certainty equivalence

ABSTRACT

For linear quadratic Gaussian problems, this paper uses two risk-sensitivity operators defined by Hansen and Sargent (2007b) to construct decision rules that are robust to misspecifications of (1) transition dynamics for state variables and (2) a probability density over hidden states induced by Bayes' law. Duality of risk sensitivity to the multiplier version of min–max expected utility theory of Hansen and Sargent (2001) allows us to compute risk-sensitivity operators by solving two-player zero-sum games. Because the approximating model is a Gaussian probability density over sequences of signals and states, we can exploit a modified certainty equivalence principle to solve four games that differ in continuation value functions and discounting of time t increments to entropy. The different games express different dimensions of concerns about robustness. All four games give rise to time consistent worst-case distributions for observed signals. But in Games I–III, the minimizing players' worst-case densities over hidden states are time inconsistent, while Game IV is an LQG version of a game of Hansen and Sargent (2005) that builds in time consistency. We show how detection error probabilities can be used to calibrate the risk-sensitivity parameters that govern fear of model misspecification in hidden Markov models.

1. Introduction

Hand in hand with advances in computer technology, hidden Markov models (HMMs) have become an essential tool for attacking dynamic estimation and decision problems (see Elliott et al., 1995; Cappé et al., 2005; Sclove, 1983; Hamilton, 1989; Sims and Zha, 2006). Hidden states can include unknown parameters, dummy variables indexing different models, hidden information variables, capital stocks, and effort levels. This paper constructs robust decision rules and estimators for a tractable class of HMM's, namely, linear-quadratic-Gaussian (LQG) Markov discounted dynamic programming problems with hidden state variables.[1] To build robust decision rules, we solve four two-player zero-sum games that, by using different continuation valuation functions exponentially to twist probability densities, focus concerns about misspecification on different aspects of a decision maker's stochastic model. The different games focus attention on different dimensions of potential misspecifications. A minimizing player helps a maximizing player design decision rules that satisfy bounds on the value of an objective function over a set of stochastic models that surround a baseline approximating model. The four games are linear-quadratic versions of games with more general functional forms

* Corresponding author at: New York University, 19 W. 4th Street, 6FL, New York, NY 10012, USA. Tel.: +1 212 998 8900; fax: +1 212 995 4186.
E-mail addresses: l-hansen@uchicago.edu (L.P. Hansen), ricardom@uchicago.edu (R. Mayer), ts43@nyu.edu (T. Sargent).

[1] See Cogley et al. (2008) for a study of robustness in a non-LQG HMM in which the hidden state indexes models unknown to a monetary authority. See Kasa (2001, 2002) for other applications of robustness to dynamic economic settings.

0165-1889/$ - see front matter
doi:10.1016/j.jedc.2010.05.004

described by Hansen and Sargent (2005, 2007b), with Game IV bringing discounting into a game analyzed by Whittle (1990) and Başar and Bernhard (1995). The LQG setting facilitates computation.

To set the stage, Section 2 describes the relationship between two classic problems without fear of model misspecification, namely, (a) a linear-quadratic-Gaussian discounted dynamic programming problem with fully observed state and (b) a linear-quadratic-Gaussian discounted dynamic programming problem with partially unobserved state. Without fear of model misspecification, a certainty equivalence principle allows one to separate the problem of estimating the hidden state from the problem of making decisions conditioned on the distribution of the state. Certainty equivalence asserts that the optimal decision is a function of the conditional mean of the hidden state and that this function is independent of the conditional volatilities in the transition equation for the hidden states as well as of the variance of the hidden state about its conditional mean. This version of certainty equivalence does not hold when we introduce fear of model misspecification, but another very useful one does. Section 3 sets out LQG problems with partially hidden states in which the decision maker fears misspecification of either (a) the distribution of stochastic shocks w^* to signals and the state transition dynamics conditioned on the entire state, or (b) the distribution of the hidden state z conditional on a history of observed signals under his approximating model, or (c) both. Sections 4–6 describe three games in which a minimizing player helps a maximizing player design a decision rule that is robust to perturbations to the distributions of w^* and z under the assumption that a time t minimizing player can disregard past perturbations of the distribution of the hidden state z. Robustness of decision rules is always relative to a value function. Altering the value function induces concerns about misspecification of different aspects of the stochastic specification and accordingly leads to different zero-sum games, as we describe in detail in Section 3.1. Game I computes a decision rule that is robust when a value function depends on both hidden and observed components of the state by solving an LQG version of recursions (20) and (21) of Hansen and Sargent (2007b). Game II computes a decision rule that is robust with respect to a value function that depends only on the observed state and the distribution of the hidden state by solving an LQG version of recursion (23) of Hansen and Sargent (2007b). Game III computes a decision rule that is robust when current returns do not depend on the hidden state by solving an LQG version of the recursion in Section 5.3 of Hansen and Sargent (2007b). Section 7 measures the time inconsistency of the worst-case distribution over hidden states that emerges in Games I and II. In Section 8, we analyze a Game IV that, like one analyzed by Whittle (1990) and Başar and Bernhard (1995), commits the decision maker to honor past distortions to distributions of hidden states. Key to attaining time consistency of distributions over hidden states is that Game IV does not discount time t contributions to entropy, while Games I–III do. Section 9 describes how to calibrate the parameter θ_2 that governs the decision maker's concerns about misspecification of the distribution of the hidden state vector by extending the procedures based on detection error probabilities that Anderson et al. (2003) and Hansen and Sargent (2008, Chapter 9) used to calibrate θ_1 in models with fully observed state vectors. Section 10 offers concluding remarks. Appendix A states a useful certainty equivalence result. Appendix B gives an alternative formulation of a robust filter under commitment. Appendix C describes a suite of Matlab programs that solve the four Games.[2]

2. Two benchmark problems

We state two classic optimization problems under full trust in a dynamic stochastic model. In the first, the decision maker observes the complete state. In the second, part of the state is hidden, impelling the decision maker to estimate it. The key finding is that with part of the state hidden, the estimation and control problems can be solved separately: first estimate the state optimally, then substitute the estimate for the observed state into the decision rule derived under the assumption that the state is observed. As we shall see in later sections, this convenient separation result under complete trust in the model does not hold when the decision maker distrusts the model.

2.1. State fully observed, model trusted

Problem 2.1. The state vector is $x_t = \begin{bmatrix} y_t \\ z_t \end{bmatrix}$ and $\begin{bmatrix} Q & P \\ P' & R \end{bmatrix}$ is a positive semi-definite matrix. Both y_t and z_t are observed at t. A decision maker chooses a state-contingent sequence of actions $\{a_t\}_{t=0}^{\infty}$ to maximize

$$-\frac{1}{2}E_0 \sum_{t=0}^{\infty} \beta^t \begin{bmatrix} a_t \\ x_t \end{bmatrix}' \begin{bmatrix} Q & P \\ P' & R \end{bmatrix} \begin{bmatrix} a_t \\ x_t \end{bmatrix} \tag{1}$$

subject to the law of motion

$$y_{t+1} = A_{11} y_t + A_{12} z_t + B_1 a_t + C_1 w_{t+1},$$

$$z_{t+1} = A_{21} y_t + A_{22} z_t + B_2 a_t + C_2 w_{t+1}, \tag{2}$$

[2] The reader who prefers to write his or her own programs and who is familiar with the deterministic discounted optimal linear regulator problem presented, for example, in Hansen and Sargent (2008, Chapter 4), will recognize how the optimal linear regulator can readily be tricked into solving games II and III.

where w_{t+1} is an iid random vector distributed as $\mathcal{N}(0,I)$, a_t is a vector of actions, and E_0 is a mathematical expectation conditioned on known initial conditions (y_0, z_0).

Guess a quadratic optimal value function

$$V(y,z) = -\frac{1}{2}\begin{bmatrix} y \\ z \end{bmatrix}' \Omega \begin{bmatrix} y \\ z \end{bmatrix} - \omega. \tag{3}$$

Let * denote next period values for variables and matrices. The Bellman equation for Problem 2.1 is

$$-\frac{1}{2}\begin{bmatrix} y \\ z \end{bmatrix}' \Omega \begin{bmatrix} y \\ z \end{bmatrix} - \omega = \max_a \left\{ -\frac{1}{2}\begin{bmatrix} a \\ y \\ z \end{bmatrix}' \begin{bmatrix} Q & P_1 & P_2 \\ P_1' & R_{11} & R_{12} \\ P_2' & R_{21} & R_{22} \end{bmatrix} \begin{bmatrix} a \\ y \\ z \end{bmatrix} - E\beta \frac{1}{2}\begin{bmatrix} y^* \\ z^* \end{bmatrix}' \Omega^* \begin{bmatrix} y^* \\ z^* \end{bmatrix} - \beta\omega^* \right\}, \tag{4}$$

where the maximization is subject to

$$y^* = A_{11}y + A_{12}z + B_1 a + C_1 w^*,$$

$$z^* = A_{21}y + A_{22}z + B_2 a + C_2 w^*, \tag{5}$$

and the mathematical expectation E is evaluated with respect to $w^* \sim \mathcal{N}(0,I)$.

Proposition 2.2. *The Bellman Eq.* (4) *induces mappings from Ω^* to Ω and from (ω^*,Ω^*) to ω. The mapping from Ω^* to Ω is a matrix Riccati difference equation that converges to a unique positive semi-definite matrix $\overline{\Omega}$ starting from any initial matrix Ω_0. The fixed point $\overline{\Omega}$ is a matrix that is a function of β, A, B, Q, P, R but is independent of the volatility matrix C that governs the 'noise statistics', i.e., the variance of x^* conditional on x. Problem* (4) *and* (5) *is an ordinary stochastic discounted optimal linear regulator problem with solution $a = -\overline{F}\left[\begin{smallmatrix} y \\ z \end{smallmatrix}\right]$, where $\overline{F}$ is independent of the volatility matrix C. The constant ω depends on C as well as on the other parameters of the problem. That $\overline{\Omega}$ and $\overline{F}$ are independent of the volatility matrix C is a manifestation of a* certainty equivalence principle (*see* Hansen and Sargent, 2008, p. 29).

2.2. State partially unobserved, model trusted

The next problem enables us to state a classic certainty equivalence result about how estimation and decision separate into two parts: (1) estimation of the hidden components of the state is done recursively with a Kalman filter that depends on C but is independent of the objective function parameters β,Q,R,P and (2) a is chosen via a *deterministic* optimal linear regulator problem that yields a linear decision rule with coefficients $\overline{F}$ that do not depend on the volatility matrix C.

Problem 2.3. A decision maker observes y_t, does not observe z_t, has a prior distribution $z_0 \sim \mathcal{N}(\check{z}_0,\Delta_0)$, and observes a sequence of signals $\{s_{t+1}\}$ whose time $t+1$ component is

$$s_{t+1} = D_1 y_t + D_2 z_t + H a_t + G w_{t+1}. \tag{6}$$

This and the following two equations constitute LQG specializations of Eqs. (1)–(3) of Hansen and Sargent (2007b):

$$y_{t+1} = \Pi_s s_{t+1} + \Pi_y y_t + \Pi_a a_t,$$

$$z_{t+1} = A_{21} y_t + A_{22} z_t + B_2 a_t + C_2 w_{t+1},$$

where $w_{t+1} \sim \mathcal{N}(0,I)$ is an i.i.d. Gaussian vector process. Substituting the signal (6) into the above equation for y_{t+1}, we obtain

$$y_{t+1} = (\Pi_s D_1 + \Pi_y) y_t + \Pi_s D_2 z_t + (\Pi_s H + \Pi_a) a_t + \Pi_s G w_{t+1},$$

which gives the y-rows in the following state-space system:

$$y_{t+1} = A_{11} y_t + A_{12} z_t + B_1 a_t + C_1 w_{t+1},$$

$$z_{t+1} = A_{21} y_t + A_{22} z_t + B_2 a_t + C_2 w_{t+1},$$

$$s_{t+1} = D_1 y_t + D_2 z_t + H a_t + G w_{t+1}, \tag{7}$$

where

$$A_{11} \doteq \Pi_s D_1 + \Pi_y, \quad A_{12} \doteq \Pi_s D_2, \quad B_1 \doteq \Pi_s H + \Pi_a, \quad C_1 \doteq \Pi_s G. \tag{8}$$

By applying Bayes' law, the decision maker constructs a sequence of posterior distributions $z_t \sim \mathcal{N}(\check{z}_t,\Delta_t), t \geq 1$, where $\check{z}_t = E[z_t | y_t, \ldots, y_1]$ for $t \geq 1$, $\Delta_t = E(z_t - \check{z}_t)(z_t - \check{z}_t)'$, and $q_t = (\check{z}_t, \Delta_t)$ is a list of sufficient statistics for the history of signals that can be expressed recursively in terms of the $(\check{z},\Delta)$ components of the following linear system:

$$y_{t+1} = A_{11} y_t + A_{12} \check{z}_t + B_1 a_t + C_1 w_{t+1} + A_{12}(z_t - \check{z}_t),$$
$$\check{z}_{t+1} = A_{21} y_t + A_{22} \check{z}_t + B_2 a_t + K_2(\Delta_t) G w_{t+1} + K_2(\Delta_t) D_2 (z_t - \check{z}_t),$$
$$\Delta_{t+1} = \mathcal{C}(\Delta_t), \tag{9}$$

where $K_2(\Delta)$ and $\mathcal{C}(\Delta)$ can be computed recursively using the Kalman filtering equations

$$K_2(\Delta) = (A_{22}\Delta D_2' + C_2 G')(D_2 \Delta D_2' + GG')^{-1}, \tag{10}$$

$$\mathcal{C}(\Delta) \equiv A_{22}\Delta A_{22}' + C_2 C_2' - K_2(\Delta)(A_{22}\Delta D_2' + C_2 G')'. \tag{11}$$

The decision maker's objective is the same as in Problem 2.1, except that his information set is now reduced to $(y_t, \check{z}_t, \Delta_t)$ at t. The current period contribution to the decision maker's objective

$$U(y,z,a) = -\left(\frac{1}{2}\right)\begin{bmatrix} a \\ y \\ z \end{bmatrix}' \begin{bmatrix} Q & P_1 & P_2 \\ P_1' & R_{11} & R_{12} \\ P_2' & R_{21} & R_{22} \end{bmatrix} \begin{bmatrix} a \\ y \\ z \end{bmatrix}$$

can be expressed as

$$\tilde{U}(y,\check{z},z-\check{z},a) = -\left(\frac{1}{2}\right)\left\{ \begin{bmatrix} a \\ y \\ \check{z} \end{bmatrix}' \begin{bmatrix} Q & P_1 & P_2 \\ P_1' & R_{11} & R_{12} \\ P_2' & R_{21} & R_{22} \end{bmatrix} \begin{bmatrix} a \\ y \\ \check{z} \end{bmatrix} + (z-\check{z})' R_{22}(z-\check{z}) + 2(z-\check{z})'(P_2 a + R_{21} y + R_{22}\check{z}) \right\},$$

whose expectation conditioned on current information $(y,\check{z},\Delta)$ equals

$$-\left(\frac{1}{2}\right)\left\{ \begin{bmatrix} a \\ y \\ \check{z} \end{bmatrix}' \begin{bmatrix} Q & P_1 & P_2 \\ P_1' & R_{11} & R_{12} \\ P_2' & R_{21} & R_{22} \end{bmatrix} \begin{bmatrix} a \\ y \\ \check{z} \end{bmatrix} + \text{trace}(R_{22}\Delta) \right\}. \tag{12}$$

Guess that the value function is

$$V(y,\check{z},\Delta) = -\frac{1}{2}\begin{bmatrix} y \\ \check{z} \end{bmatrix}' \Omega \begin{bmatrix} y \\ \check{z} \end{bmatrix} - \omega \tag{13}$$

and choose Ω and ω to verify the Bellman equation

$$-\frac{1}{2}\begin{bmatrix} y \\ \check{z} \end{bmatrix}' \Omega \begin{bmatrix} y \\ \check{z} \end{bmatrix} - \omega = \max_a E\left\{ -\frac{1}{2}\begin{bmatrix} a \\ y \\ \check{z} \end{bmatrix}' \begin{bmatrix} Q & P_1 & P_2 \\ P_1' & R_{11} & R_{12} \\ P_2' & R_{21} & R_{22} \end{bmatrix} \begin{bmatrix} a \\ y \\ \check{z} \end{bmatrix} - \frac{1}{2}\text{trace}(R_{22}\Delta) - E\beta\frac{1}{2}\begin{bmatrix} y^* \\ \check{z}^* \end{bmatrix}' \Omega^* \begin{bmatrix} y^* \\ \check{z}^* \end{bmatrix} - \beta\omega^* \right\}, \tag{14}$$

where the maximization is subject to the *innovation representation*

$$y^* = A_{11}y + A_{12}\check{z} + B_1 a + \{C_1 w^* + A_{12}(z-\check{z})\},$$

$$\check{z}^* = A_{21}y + A_{22}\check{z} + B_2 a + \{K_2(\Delta)Gw^* + K_2(\Delta)D_2(z-\check{z})\},$$

$$\Delta^* = \mathcal{C}(\Delta), \tag{15}$$

with given initial conditions $(y_0, \check{z}_0)$ and where the mathematical expectation E is evaluated with respect to $w^* \sim \mathcal{N}(0,I)$ and $z-\check{z} \sim \mathcal{N}(0,\Delta)$. Notice that the matrices A_{11}, A_{12}, B_1, A_{21}, A_{22}, B_2 defining the systematic parts of the laws of motion (5) and (15), i.e., the parts other than the linear combinations of the shocks w^*, $(z-\check{z})$, are identical. In light of the objective functions in the two problems, this fact implies:

Proposition 2.4. *Associated with Bellman Eq.* (14) *is the same matrix Riccati difference equation mapping Ω^* into Ω that characterized Problem* 2.1. *It converges to a unique positive semi-definite solution $\overline{\Omega}$ starting from any initial value Ω_0. The fixed point $\overline{\Omega}$ is a function of β,A,B,Q,P,R but is independent of the matrices C_1,C_2,K_2,D_2,Δ that determine the volatilities of y^*, $\check{z}^*$ conditional on $y,\check{z}$. Because the Riccati equation for Ω is identical with the one associated with Problem* 2.1, *the fixed point $\overline{\Omega}$ and the matrix $\overline{F}$ in the decision rule $a = -\overline{F}\left[\begin{smallmatrix} y \\ \check{z} \end{smallmatrix}\right]$ are identical with their counterparts in Problem* 2.2. *These outcomes justify a separation of optimization and estimation that embody a certainty-equivalence principle. The sequence of constants $\{\omega_t\}$ depends on the sequences $\{\Delta_t\}$,* trace$(R_{22}\Delta_t)$ *and differs from its counterpart in Problem* 2.1.

Remark 2.5. The matrix $\overline{\Omega}$ in the quadratic form in the optimal value function and the matrix $\overline{F}$ in the decision rule for Problems 2.1 and 2.3 can be computed by solving the same *deterministic optimal linear regulator problem* (see Hansen and Sargent, 2008, Chapter 4) that we can construct by setting to zero the volatility matrices multiplying shocks in the respective problems. After that, the constants ω can be computed by solving appropriate versions of the usual recursion mapping ω^* and other objects into ω.

3. State partially unobserved, model *dis* trusted

We modify Problem 2.3 by positing a decision maker who distrusts the joint distribution for $\{y_t, z_t\}_{t=0}^{\infty}$ that is implied by systems (9)–(11) and therefore wants a robust decision rule, i.e., one that attains a value that is guaranteed to exceed a bound over a set of perturbed distributions. To formulate the problem recursively, we express his distrust in terms of two types of conditional distributions that are components of his *approximating model* (15):

1. The distribution of w^* conditional either on a complete information set $(y,z,\check{z},\Delta)$ or on the incomplete information set $(y,\check{z},\Delta)$.
2. The distribution of $(z-\check{z})$ conditional on the history of signals.

Following Hansen and Sargent (2005, 2007b), we compute robust decision rules by replacing the expectation operator in Bellman Eq. (14) with compositions of two risk-sensitivity operators, one of which adjusts continuation values for possible misspecification of the conditional densities of w^*, the other of which adjusts for possible misspecfication of $z-\check{z}$. We exploit the insight that application of the two risk-sensitivity operators yield indirect utility functions of malevolent players who choose distributions of w^* and $z-\check{z}$, respectively, to minimize the objective of the maximizing player.[3] By responding to the minimizing choices of probabilities, the maximizing player constructs a decision rule that is robust to perturbations to the distributions of w^* and $z-\check{z}$.

3.1. Robustness is context specific

A robust decision rule is *context specific* in the sense that it depends on the preference parameters in $(\beta,U(y,z,a))$ and also on details of the stochastic perturbations to the approximating model (15) that concern the decision maker.

We create alternative two-player zero-sum games that differ in either the one-period objective function $U(y,z,a)$ or the stochastic perturbations to the approximating model, in particular, the conditioning sets for the densities of w^* and $z-\check{z}$. Each game expresses concerns about robustness in terms of two positive penalty parameters, a θ_1 that measures the decision maker's distrust of the distribution of w^*, and a θ_2 that measures the decision maker's distrust of the distribution of $z-\check{z}$ that emerges from applying Bayes' law using the approximating model.

In these games, the ordinary version of certainty equivalence does not prevail: decision rules now depend vitally on matrices governing conditional volatilities. However, a modified version of certainty equivalence described by Hansen and Sargent (2008, p. 33) and Appendix A does apply. It allows us to compute robust decisions by solving deterministic two-player zero-sum games while keeping track only of the distorted *means* of perturbed distributions together with conditional volatility matrices associated with the approximating model.

Different zero-sum two-player games focus a decision maker's distrust on different aspects of the baseline stochastic model (15). Hansen and Sargent (2007b) used Games I–III to generate Bellman equations that closely resemble (14). These games acquire convenient recursive structures by accepting time inconsistency in equilibrium worst-case distributions for the hidden state, as we emphasize in Section 7. A linear-quadratic version of a game proposed by Hansen and Sargent (2005), called Game IV, is different and builds in time consistency of those distributions by not discounting time t contributions to entropy and by making the minimizing player choose once and for all at time 0.

The four Games differ in timing protocols and the information ascribed to the minimizing player who, by distorting probability distributions, helps the maximizing player achieve robustness. In Games I–III, there are sequences of minimizing players.

- Game I (an LQG version of recursions (20) and (21) of Hansen and Sargent, 2007b) starts with a date $t+1$ value function that depends on $y_{t+1},z_{t+1},\check{z}_{t+1}$. Then a minimizing player at $t\geq 0$ distorts the distribution of w_{t+1} conditional on $y_t,z_t,\check{z}_t$, as restrained by a positive penalty parameter θ_1. This leads to a date t value function that conditions on $y_t,z_t,\check{z}_t$. Then a minimizing player distorts the distribution of $z_t-\check{z}_t$ conditional on $y_t,\check{z}_t$, as restrained by a positive penalty parameter θ_2.
- Game II (an LQG version of recursions (23) of Hansen and Sargent, 2007b) starts with a date $t+1$ value function that depends on $y_{t+1},\check{z}_{t+1}$. A minimizing agent distorts the distribution of w_{t+1} conditioned on $y_t,z_t,\check{z}_t$, as restrained by a positive penalty parameter θ_1. This leads to a date t value function that depends on $y_t,z_t,\check{z}_t$. Then a minimizing player distorts the distribution of $z_t-\check{z}_t$ conditional on $y_t,\check{z}_t$, as restrained by a positive penalty parameter θ_2. A date t value function conditions on y_t and $\check{z}_t$.
- Game III (an LQG version of the recursion in Hansen and Sargent, 2007b, Section 5.3) is a special case of Game II in which the decision maker's one-period objective function does not depend on z_t and in which $\theta_1=\theta_2$.

The arguments of their value functions distinguish Games I and II.

Game IV (an LQG version of the game with undiscounted entropy in Hansen and Sargent, 2005) has a single minimizing player who chooses once and for all at time 0.

- In Game IV, a time 0 decision maker observes $y_t,\check{z}_t$ at time t and chooses distortions of the distribution of w_{t+1} conditional on the history of y_s, z_s for $s=0,\ldots,t$, as well as a time 0 distortion to the distribution of $z_0-\check{z}_0$.

[3] This is the insight that connects robust control theory to risk-sensitivity. See Hansen and Sargent (2008). See Cerreia et al. (2008) for a general representation of uncertainty averse preferences in terms of indirect utility functions for a minimizing player who chooses probabilities.

Each of the four games implies worst-case distortions to the mean vector and covariance matrix of the shocks w_{t+1} in (9) and to the mean and covariance $(\check{z}_t,\Delta_t)$ that emerge from the Kalman filter. The worst-case means feed back on the state in ways that help the decision maker design robust decisions.

4. Game I: value function depends on $(y,z,\check{z},\Delta)$

This game corresponds to an LQG version of Hansen and Sargent (2007b, recursions (20) and (21)) and lets the minimizing player but not maximizing player observe z. That information advantage induces the decision maker to explore the fragility of his decision rule with respect to misspecifications of the dynamics conditional on the entire state. Because the maximizing and minimizing players have different information sets, we solve this game in two steps. The first step chooses a distortion $\tilde{v}$ to the mean of w^* and a distortion u to the mean of $z-\check{z}$ conditioned on information $(y,\check{z},\Delta)$ available to the maximizing player. The second step chooses a distortion v to the mean of w^* conditioned on the larger information set $(y,\check{z},\Delta,z)$ available to the minimizing player.

Step 1 (Choosing $\tilde{v}$ and u conditioned on $(y,\check{z},\Delta)$): This step corresponds to solving problem (21) of Hansen and Sargent (2007b). Let $W(y^*,\check{z}^*,\Delta^*,z^*)$ be a quadratic function of next period's state variables. In terms of the state variables $(y,\check{z},\Delta)$, the law of motion for $(y,z,\check{z},\Delta)$ can be written as

$$\begin{aligned} y^* &= A_{11}y+A_{12}\check{z}+B_1a+C_1w^*+A_{12}(z-\check{z}),\\ z^* &= A_{21}y+A_{22}\check{z}+B_2a+C_2w^*+A_{22}(z-\check{z}),\\ \check{z}^* &= A_{21}y+A_{22}\check{z}+B_2a+K_2(\Delta)Gw^*+K_2(\Delta)D_2(z-\check{z}),\\ \Delta^* &= \mathcal{C}(\Delta), \end{aligned} \tag{16}$$

where $w^*\sim\mathcal{N}(0,I)$ and $z-\check{z}\sim\mathcal{N}(0,\Delta)$. We replace these distributions with the distorted distributions $w^*\sim\mathcal{N}(\tilde{v},\Sigma)$ and $z-\check{z}\sim\mathcal{N}(u,\Gamma)$. By feeding back on prior states, $\tilde{v}$ and u can represent possibly misspecified dynamics in the approximating model. At this point, we use a modified certainty equivalence result to form a law of motion for a deterministic two-player zero-sum game that will yield a decision rule that solves the stochastic two-player zero-sum game (21) of Hansen and Sargent (2007b) that interests us. We replace w^* with the distorted mean vector $\tilde{v}$ and $z-\check{z}$ with the distorted mean vector u. The modified certainty equivalence principle in Hansen and Sargent (2008, p. 33) and Appendix A asserts that we can solve (21) of Hansen and Sargent (2007b) by replacing it with a deterministic two-player zero-sum game that treats the distorted means $\tilde{v}$ and u as variables under the control of a minimizing player. Omitted stochastic terms affect constants in value functions, but not decision rules. Replacing shocks with distorted means gives us

$$\begin{aligned} y^* &= A_{11}y+A_{12}\check{z}+B_1a+C_1\tilde{v}+A_{12}u,\\ z^* &= A_{21}y+A_{22}\check{z}+B_2a+C_2\tilde{v}+A_{22}u,\\ \check{z}^* &= A_{21}y+A_{22}\check{z}+B_2a+K_2(\Delta)G\tilde{v}+K_2(\Delta)D_2u,\\ \Delta^* &= \mathcal{C}(\Delta). \end{aligned} \tag{17}$$

Problem 4.1. Let θ_1 and θ_2 be positive scalars. For a quadratic value function $W(y,\check{z},\Delta,z)$, to be computed in step 2, choose an action a and accompanying distorted mean vectors u, $\tilde{v}$ by solving Bellman equation

$$\max_a \min_u \left[\tilde{U}(y,\check{z},z-\check{z},a)+\theta_2\frac{u'\Delta^{-1}u}{2}+\min_{\tilde{v}}\left(\beta W(y^*,\check{z}^*,\Delta^*,z^*)+\theta_1\frac{\tilde{v}'\tilde{v}}{2}\right)\right], \tag{18}$$

where the optimization is subject to the law of motion (17). Minimization over $\tilde{v}$ implements risk-sensitivity operator T^1 and minimization over u implements T^2 in the stochastic problem (21) of Hansen and Sargent (2007b). A robust decision rule attains the right side of (18) and takes the form $a=-F(\Delta)\left[\begin{smallmatrix} y\\ \check{z}\end{smallmatrix}\right]$. To make the extremization on the right side of Bellman Eq. (18) well posed, (θ_1,θ_2) must be large enough that

$$\begin{bmatrix} \theta_2\Delta^{-1}-R_{22} & 0\\ 0 & \theta_1 I\end{bmatrix}-\beta\begin{bmatrix} A_{12} & C_1\\ A_{22} & C_2\\ K_2(\Delta)D_2 & K_2(\Delta)G\end{bmatrix}'\Omega^*(\Delta^*)\begin{bmatrix} A_{12} & C_1\\ A_{22} & C_2\\ K_2(\Delta)D_2 & K_2(\Delta)G\end{bmatrix}$$

is positive definite.

Remark 4.2. The matrices $C_1,C_2,A_{12},A_{22},K_2(\Delta)D_2$ that determine conditional volatilities in the approximating model (15) influence the maximizing player's choice of a because they determine the minimizing player's decisions $\tilde{v}$, u and therefore the future state.

Step 2 (Choosing v with continuation value conditioned on $(y,\check{z},\Delta,z)$): This step constructs a continuation value function $W(y,\check{z},\Delta,z)$ by allowing the minimizing player to condition on z as well as on $(y,\check{z},\Delta)$. This corresponds to solving (20) of Hansen and Sargent (2007b). To facilitate conditioning on z, rewrite the law of motion as

$$\begin{bmatrix} y^* \\ z^* \\ \check{z}^* \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} & 0 \\ A_{21} & A_{22} & 0 \\ A_{21} & K_2(\Delta)D_2 & A_{22}-K_2(\Delta)D_2 \end{bmatrix} \begin{bmatrix} y \\ z \\ \check{z} \end{bmatrix} - \begin{bmatrix} B_1 \\ B_2 \\ B_2 \end{bmatrix} F(\Delta) \begin{bmatrix} y \\ \check{z} \end{bmatrix} + \begin{bmatrix} C_1 \\ C_2 \\ K_2(\Delta)G \end{bmatrix} v \tag{19}$$

together with

$$\Delta^* = \mathcal{C}(\Delta). \tag{20}$$

Here v is the distorted mean of w^* conditioned on $(y,z,\check{z},\Delta)$, while $\tilde{v}$ in step 1 is the distorted mean of w^* conditional on $(y,\check{z},\Delta)$.

Problem 4.3. Posit a quadratic value function

$$W(y,\check{z},\Delta,z) = -\frac{1}{2} \begin{bmatrix} y \\ z \\ \check{z} \end{bmatrix}' \Omega(\Delta) \begin{bmatrix} y \\ z \\ \check{z} \end{bmatrix} - \omega$$

and update it via

$$W(y,\check{z},\Delta,z) = U(y,z,a) + \min_{v} \left\{ \beta W^*(y^*,\check{z}^*,\Delta^*,z^*) + \theta_1 \frac{v'v}{2} \right\}, \tag{21}$$

where the minimization is subject to the law of motion (19) and (20). For the minimization problem on the right side of Bellman Eq. (21) to be well posed, we require that θ_1 be large enough that $\theta_1 I - \beta \overline{C}(\Delta)'[\Omega^* \circ \mathcal{C}(\Delta)]\overline{C}(\Delta)$ is positive definite, where

$$\overline{C}(\Delta) = \begin{bmatrix} C_1 \\ C_2 \\ K_2(\Delta)G \end{bmatrix}.$$

Remark 4.4. We use the modified certainty-equivalence principle described by Hansen and Sargent (2008, Chapter 2) and Appendix A. After we compute the worst-case conditional means, v, u, it is easy to compute the corresponding worst-case conditional variances $\Sigma(\Delta)$, $\Gamma(\Delta)$ of w^* and $z-\check{z}$, respectively, as

$$\Sigma(\Delta) \doteq \left(I - \frac{\beta}{\theta_1} \overline{C}(\Delta)'[\Omega^* \circ \mathcal{C}(\Delta)]\overline{C}(\Delta) \right)^{-1} \tag{22}$$

and

$$\Gamma(\Delta) \doteq \left(\Delta^{-1} - \frac{1}{\theta_2} R_{22} - \frac{1}{\theta_2} [0 \;\; I \;\; 0] \Omega(\Delta) \begin{bmatrix} 0 \\ I \\ 0 \end{bmatrix} \right)^{-1}. \tag{23}$$

These formulas are valid when the matrices on the right side of (22) and (23) are positive definite.

Remark 4.5. The decision rule $a = -F(\Delta)\left[\begin{smallmatrix} y \\ \check{z} \end{smallmatrix}\right]$ that solves the infinite horizon problem also solves a stochastic counterpart that is formulated in terms of the T^1 and T^2 risk-sensitivity operators (Hansen and Sargent, 2007b, Eq. (21)).

5. Game II: value function depends on $(\mathbf{y},\check{\mathbf{z}},\Delta)$

This game withdraws the Game I information advantage from the minimizing player and works with a transition law for only $(y,\check{z})$. The game solves an LQG version of Hansen and Sargent (2007b, recursion (23)). To exploit the modified certainty equivalence principle of Appendix A, we replace w^* with $\tilde{v}$ and $(z-\check{z})$ with u in the stochastic law of motion (15) to obtain

$$\begin{bmatrix} y^* \\ \check{z}^* \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \begin{bmatrix} y \\ \check{z} \end{bmatrix} + \begin{bmatrix} B_1 \\ B_2 \end{bmatrix} a + \begin{bmatrix} C_1 \\ K_2(\Delta)G \end{bmatrix} \tilde{v} + \begin{bmatrix} A_{12} \\ K_2(\Delta)D_2 \end{bmatrix} u. \tag{24}$$

Problem 5.1. Guess a quadratic value function

$$V(y,\check{z}) = -\frac{1}{2}\begin{bmatrix} y \\ \check{z} \end{bmatrix}' \Omega(\Delta) \begin{bmatrix} y \\ \check{z} \end{bmatrix} - \omega$$

and form the Bellman equation

$$V(y,\check{z}) = \max_{a} \min_{u,\tilde{v}} \left[\tilde{U}(y,\check{z},z-\check{z},a) + \theta_2 \frac{u'\Delta^{-1}u}{2} + \beta V(y^*,\check{z}^*) + \theta_1 \frac{\tilde{v}'\tilde{v}}{2} \right], \tag{25}$$

where the optimization is subject to the law of motion (24). This can be formulated as a *deterministic optimal linear regulator* problem. For the extremization problem on the right side of the Bellman equation to be well posed, we require that (θ_1,θ_2) be large enough that

$$\begin{bmatrix} \theta_2 \Delta^{-1} - R_{22} & 0 \\ 0 & \theta_1 I \end{bmatrix} - \beta \begin{bmatrix} A_{12}' & D_2' K_2(\Delta)' \\ C_1' & G' K_2(\Delta)' \end{bmatrix} \Omega(\Delta^*) \begin{bmatrix} A_{12} & C_1 \\ K_2(\Delta) D_2 & K_2(\Delta) G \end{bmatrix}$$

is positive definite. After Bellman Eq. (25) has been solved, worst-case variances $\Sigma(\Delta)$ and $\Gamma(\Delta)$ of w^* and $z-\check{z}$, respectively, can be computed using standard formulas.

6. Game III: $\theta_1 = \theta_2$ and no hidden states in objective

Game III solves an LQG version of the recursion described in Hansen and Sargent (2007b, Section 5.3). Game III is a special case of Game II that features situations in which[4]:

1. The current period objective function depends on (y, a) but not on z.
2. As in Game II, the decision maker and the minimizing player both have access to the reduced information set $(y,\check{z},\Delta)$.
3. The multipliers $\theta_1 = \theta_2 = \theta$.

The one period objective is

$$\hat{U}(y,a) = -\frac{1}{2}\begin{bmatrix} a \\ y \end{bmatrix}' \begin{bmatrix} Q & P \\ P' & R \end{bmatrix} \begin{bmatrix} a \\ y \end{bmatrix}.$$

The law of motion for the stochastic system is

$$\begin{bmatrix} y^* \\ \check{z}^* \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \begin{bmatrix} y \\ \check{z} \end{bmatrix} + \begin{bmatrix} B_1 \\ B_2 \end{bmatrix} a + \left\{ \begin{bmatrix} C_1 \\ K_2(\Delta) G \end{bmatrix} w^* + \begin{bmatrix} A_{12} \\ K_2(\Delta) D_2 \end{bmatrix} (z - \check{z}) \right\}.$$

Using a Cholesky decomposition of the covariance matrix of the composite shock in braces, we represent it in terms of a new normalized composite shock $\tilde{w}$ as

$$\begin{bmatrix} y^* \\ \check{z}^* \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \begin{bmatrix} y \\ \check{z} \end{bmatrix} + \begin{bmatrix} B_1 \\ B_2 \end{bmatrix} a + \tilde{C} \tilde{w}^*,$$

where $\tilde{w}^* \sim \mathcal{N}(0,I)$ and

$$\tilde{C}\tilde{w}^* = \left\{ \begin{bmatrix} C_1 \\ K_2(\Delta) G \end{bmatrix} w^* + \begin{bmatrix} A_{12} \\ K_2(\Delta) D_2 \end{bmatrix} (z - \check{z}) \right\}.$$

To form an appropriate deterministic optimal linear regulator problem, we replace $\tilde{w}^*$ with the distorted mean $\tilde{v}$ to form the law of motion

$$\begin{bmatrix} y^* \\ \check{z}^* \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \begin{bmatrix} y \\ \check{z} \end{bmatrix} + \begin{bmatrix} B_1 \\ B_2 \end{bmatrix} a + \tilde{C} \tilde{v}. \tag{26}$$

Problem 6.1. For a quadratic value function

$$W(y,\check{z}) = -\frac{1}{2}\begin{bmatrix} y \\ \check{z} \end{bmatrix}' \Omega \begin{bmatrix} y \\ \check{z} \end{bmatrix} - \omega,$$

[4] Hansen et al. (2002, Eqs. (45) and (46)) used a version of Game III to study the effects of robust filtering on asset pricing.

solve the Bellman equation

$$W(y,\check{z}) = \max_{a} \min_{\tilde{v}} \left\{ \hat{U}(y,a) + \beta W^*(y^*,\check{z}^*) + \theta \frac{\tilde{v}'\tilde{v}}{2} \right\}, \tag{27}$$

where the optimization is subject to (20) and (26). A robust decision rule is $a = -F(\Delta)[^y_{\check{z}}]$ and the worst-case mean of $\tilde{w}^*$ is $\tilde{v} = H(\Delta)[^y_{\check{z}}]$; the worst-case variance of $\tilde{w}^*$ can be found as described in Hansen and Sargent (2008, Chapter 3). For the minimization part of the problem on the right side of Bellman equation (27) to be well posed, we require that θ be large enough that

$$\theta I - \beta \tilde{C}(\Delta)'[\Omega^* \circ \mathcal{C}(\Delta)]\tilde{C}(\Delta)$$

is positive definite.

Remark 6.2. Problem 6.1 is an ordinary robust control problem with an observed state (see Hansen and Sargent, 2008, Chapter 2) that does not separately consider perturbations to the distributions of the components w^* and $z - \check{z}$ that contribute to $\tilde{v}$.

7. Time inconsistency of worst-case z distributions

While Game III focuses on possible misspecification of a composite shock and so does not distinguish between errors coming from misspecifications of the separate distributions w^* and $z - \check{z}$, Games I and II do explicitly separate misspecification of the w^* shock from the $z - \check{z}$ shock. They allow the minimizing player at time t to distort the distribution of $z - \check{z}$ anew each period and to disregard the distortions to the distribution of $z - \check{z}$ that are implied by prior distortions to the distributions of past w's and $z - \check{z}$'s.[5] This feature of Games I and II makes the worst-case distributions for z time inconsistent.

We illustrate the time inconsistency of worst-case beliefs about z in the context of Game I. For Game I, the worst-case distribution of w^* conditioned on $(y,\check{z},\Delta)$ is $\mathcal{N}(v,\Sigma(\Delta))$, where v is a linear function of $y,\check{z}$ that can be expressed as

$$v = \beta[\theta_1 I - \beta \overline{C}(\Delta)'[\Omega^* \circ \mathcal{C}(\Delta)]\overline{C}(\Delta)]^{-1} \overline{C}(\Delta)'[\Omega^* \circ \mathcal{C}(\Delta)]\overline{A} \begin{bmatrix} y \\ z \\ \check{z} \end{bmatrix}, \tag{28}$$

and $\Sigma(\Delta)$ is given by (22). The worst-case mean of z is $\check{z} + u$, where u is given by an expression of the form $u = -\tilde{F}_{21}(\Delta)y - \tilde{F}_{23}(\Delta)\check{z}$ and the worst-case covariance of z is given by (23). Using these in $z^* = A_{21}y + A_{22}z + B_2a + C_2w^*$ shows that the mean of z^* implied by the current worst-case conditional distribution is

$$\tilde{z}^* = A_{21}y + A_{22}(\check{z} + u) + B_2a + K_2(\Gamma(\Delta))(s^* - D_1y - D_2(\check{z} + u) - Ha) \tag{29}$$

and that the implied covariance of z^* is

$$\tilde{\Delta}^* = (A_{22} - K_2(\Gamma(\Delta))D_2)\Gamma(\Delta)(A_{22} - K_2(\Gamma(\Delta))D_2)' + C_2\Sigma(\Delta)C_2'. \tag{30}$$

Eqs. (29) and (30) give the mean and covariance of z^* conditioned on s^* and the history of s implied by the current worst-case model. Without commitment, these do not equal the mean and covariance, respectively, of the worst-case distribution that our decision maker synthesizes next period by adjusting the sufficient statistics $(\check{z}^*, \Delta^*)$ for the conditional distribution of z^* that emerges from the ordinary Kalman filter. Thus,

$$\tilde{z}^* \neq (\check{z}^* + u^*),$$

$$\tilde{\Delta}^* \neq \mathrm{C}[\Gamma(\Delta)]. \tag{31}$$

Gaps between the left and right sides indicate time inconsistency in the worst-case distributions of z_t.

It bears emphasizing that time inconsistency does not appear in the worst-case distribution for signals given the signal history and manifests itself only when we unbundle the distorted signal distribution into separate components coming from distortions to the w^* and $z - \check{z}$ distributions.

As stressed in Hansen and Sargent (2005), time-inconsistency of the worst-case distributions of z is a consequence of our decision to set up Games I and II by following Hansen and Sargent (2007b) in discounting time t contributions both to entropy and to utility.[6] In the next section, we follow Hansen and Sargent (2005) by not discounting entropy and positing a single minimizing player who chooses distortions once-and-for-all at time 0.

[5] For more about this feature, see Hansen and Sargent (2008, Chapter 17).

[6] Epstein and Schneider (2003, 2007) also tolerate such inconsistencies in their recursive formulation of multiple priors. As in our Game II, they work with continuation values that depend on signal histories; and they distort the distribution of the future signal conditional on the hidden states, not just the distribution conditional on the signal history.

8. Game IV: commitment to prior distortions in z distributions

Games I–III adopt sequential timing protocols that give a time t decision maker the freedom to distort afresh the distribution $\mathcal{N}(\check{z}_t, \Delta_t)$ that emerges from Bayes' law as applied to the approximating model and to disregard distortions to the distribution of z_t that are implied by the approximating model's transition law for z together with distortions to the distribution of earlier z's. We now solve a Game IV that implements a linear-quadratic version of a game of Hansen and Sargent (2005) that imposes commitment to prior distortions.

8.1. The problem

The limited information problem under commitment is[7]

$$\max_{\{a_t\}} \min_{h_0 \in \mathcal{H}_0} \min_{\{m_{t+1} \in \mathcal{M}_{t+1}\}} E \sum_{t=0}^{\infty} M_t \left(-\beta^t \left(\frac{1}{2}\right) \begin{bmatrix} a_t \\ x_t \end{bmatrix}' \begin{bmatrix} Q & P \\ P' & R \end{bmatrix} \begin{bmatrix} a_t \\ x_t \end{bmatrix} + \theta m_{t+1} \log m_{t+1} | \mathcal{S}_0 \right) + \theta E(h_0 \log h_0 | \mathcal{S}_0) \tag{32}$$

subject to

$$x_{t+1} = Ax_t + Ba_t + Cw_{t+1},$$

$$s_{t+1} = Dx_t + Ha_t + Gw_{t+1},$$

$$M_{t+1} = m_{t+1} M_t,$$

$$M_0 = h_0, \tag{33}$$

where h_0 is a non-negative random variable that is measurable with respect to y_0, z_0 and whose mean is 1; and m_{t+1} a non-negative random variable that is measurable with respect to the history of $y_s, z_s, s = t+1, t, t-1, \ldots, 0$ and whose expectation conditioned on the history up to t is 1. The decision maker's choice of h_0 at time 0 distorts the prior distribution of z_0, while his distortions of the distribution of z_t for future t's are implied by his time 0 choice of the sequence $\{m_{t+1}\}_{t=0}^{\infty}$. This captures how this game builds in commitment to prior distortions. Hansen and Sargent (2005) show that this problem can be solved in the following two steps.

8.2. Step 1: Solve a problem with observed states and without random shocks

We first solve the following two-player zero-sum game with no uncertainty. The problem is

$$\max_{\{a_t\}} \min_{\{v_t\}} \left(\frac{1}{2}\right) \sum_{t=0}^{\infty} \left(-\beta^t \begin{bmatrix} a_t \\ x_t \end{bmatrix}' \begin{bmatrix} Q & P \\ P' & R \end{bmatrix} \begin{bmatrix} a_t \\ x_t \end{bmatrix} + \theta |v_t|^2 \right) \tag{34}$$

subject to

$$x_{t+1} = Ax_t + Ba_t + Cv_t. \tag{35}$$

Notice that one-period utilities are discounted, but increments $|v_t|^2$ to entropy are not.

For sufficiently large values of θ, the Markov perfect equilibrium gives rise to a date t value function that is quadratic. Inclusive of discounting,[8] we denote it

$$-\frac{\beta^t}{2} (x_t)' \Omega_t x_t.$$

Define

$$\tilde{Q}_t \doteq \begin{bmatrix} Q & 0 \\ 0 & -\beta^{-t} \theta I \end{bmatrix},$$

[7] The distortion m_{t+1} is a likelihood ratio that changes the distribution of w_{t+1} from a normal distribution with mean zero and covariance matrix I to a normal distribution with a mean v_t that is given by the second equation of (36) and a covariance matrix Υ_t, where $\Upsilon_t^{-1} = I - (1/\theta) C' \Omega_{t+1} C \beta^{t+1}$. The distortion m_{t+1} equals

$$m_{t+1} = \exp\left[-\frac{1}{2} (w_{t+1} - v_t)' (\Upsilon_t)^{-1} (w_{t+1} - v_t) + \frac{1}{2} w_{t+1} \cdot w_{t+1} - \frac{1}{2} \log\det \Upsilon_t \right].$$

A simple calculation shows that

$$E(m_{t+1} \log m_{t+1} | \mathcal{X}_t) = \frac{1}{2} [|v_t|^2 + \mathrm{trace}(\Upsilon_t - I) - \log\det \Upsilon_t],$$

where the component terms $\frac{1}{2}|v_t|^2$ and $\mathrm{trace}(\Upsilon_t - I) - \log\det \Upsilon_t$ are both non-negative.

[8] This problem is well posed only for sufficiently large values of θ. See Başar and Bernhard (1995, Lemma 3.1).

$$\tilde{P} \doteq \begin{bmatrix} P \\ 0 \end{bmatrix},$$

$$\tilde{R} \doteq R - \tilde{P}'(\tilde{Q}_t)^{-1}\tilde{P} = R - P'Q^{-1}P,$$

$$\tilde{B} \doteq [B \;\; C],$$

$$\tilde{A} \doteq A - \tilde{B}(\tilde{Q}_t)^{-1}\tilde{P} = A - BQ^{-1}P.$$

The robust a_t and the worst-case v_t are

$$\begin{bmatrix} a_t \\ v_t \end{bmatrix} = -\begin{bmatrix} Q + B'\beta\Omega_{t+1}B & \beta B'\Omega_{t+1}C \\ \beta C'\Omega_{t+1}B & \beta C'\Omega_{t+1}C - \beta^{-t}\theta I \end{bmatrix}^{-1} \begin{bmatrix} \beta B'\Omega_{t+1}A + P \\ \beta C'\Omega_{t+1}A \end{bmatrix} x_t = -(\beta[\tilde{Q}_t + \beta\tilde{B}'\Omega_{t+1}\tilde{B}]^{-1}\tilde{B}'\Omega_{t+1}\tilde{A} + (\tilde{Q}_t)^{-1}\tilde{P})x_t, \quad (36)$$

where the matrix Ω_t in the value function satisfies the Riccati equation

$$\Omega_t = \tilde{R} + \beta\tilde{A}'\Omega_{t+1}\tilde{A} - \beta\tilde{A}'\Omega_{t+1}\tilde{B}[\tilde{Q}_t + \beta\tilde{B}'\Omega_{t+1}\tilde{B}]^{-1}\tilde{B}'\Omega_{t+1}\tilde{A}. \quad (37)$$

(Also see Başar and Bernhard, 1995, p. 272.)

When $\beta < 1$, as $t \to +\infty$, the solution for Ω_t converges to one that would be obtained under a no-robustness ($\theta = \infty$) specification, v_t converges to zero, and the limiting control law converges to that associated with $\theta = \infty$ (i.e., the one associated with no fear of model misspecification). When $\theta < +\infty$, the decision maker is concerned about robustness, but that concern diminishes over time. The dissipation of concerns about robustness with the passage of time is a direct consequence of the different discounting of one-period returns (they are discounted) and one-period entropies (they are not discounted).

8.3. *Step 2: Given $\{\Omega_t\}$, compute the filter*

Hansen and Sargent (2005) derive the following recursions for the robust estimates. Starting from the sufficient statistics $(\check{z}_0, \check{\Delta}_0)$ that describe the decision maker's prior $z_0 \sim \mathcal{N}(\check{z}_0, \check{\Delta}_0)$, for $t \geq 0$ iterate on[9]

$$\hat{z}_t = \check{z}_t + \left[(\check{\Delta}_t)^{-1} - \frac{\beta^t}{\theta}[0 \;\; I]\Omega_t \begin{bmatrix} 0 \\ I \end{bmatrix}\right]^{-1} \frac{\beta^t}{\theta}[0 \;\; I]\Omega_t \begin{bmatrix} y_t \\ \check{z}_t \end{bmatrix}, \quad (38)$$

$$a_t = -[I \;\; 0][\beta(\tilde{Q}_t + \beta\tilde{B}'\Omega_{t+1}\tilde{B})^{-1}\tilde{B}'\Omega_{t+1}\tilde{A} + (\tilde{Q}_t)^{-1}\tilde{P}] \begin{bmatrix} y_t \\ \hat{z}_t \end{bmatrix}, \quad (39)$$

$$(\tilde{\Delta}_t)^{-1} = (\check{\Delta}_t)^{-1} - \frac{\beta^t}{\theta}R_{22}, \quad (40)$$

$$\tilde{z}_t = \check{z}_t + \frac{\beta^t}{\theta}\tilde{\Delta}_t[P_2'a_t + R_{12}'y_t + R_{22}\check{z}_t], \quad (41)$$

$$\check{z}_{t+1} = \mathsf{M}(y_t, \tilde{z}_t, a_t, s_{t+1}, \tilde{\Delta}_t), \quad (42)$$

$$\check{\Delta}_{t+1} = \mathsf{C}(\tilde{\Delta}_t), \quad (43)$$

where

$$\mathsf{M}(y, \check{z}, a, s^*, \Delta) \doteq A_{21}y + A_{22}\check{z} + B_2a + K_2(\Delta)(s^* - D_1y - D_2\check{z} - Ha). \quad (44)$$

Here $\check{z}^* = \mathsf{M}(y, \check{z}, a, s^*, \Delta) \doteq A_{21}y + A_{22}\check{z} + B_2a + K_2(\Delta)(s^* - D_1y - D_2\check{z} - Ha)$ would be the update of $\check{z}$ associated with the usual Kalman filter. When hidden states appear in the one-period utility function, the commitment feature of the problem induces adjustment (40) to the estimates coming from the Kalman filter. This adjustment vanishes when the utility function contains no hidden states.[10]

Following the robust control literature (e.g., Başar and Bernhard, 1995; Whittle, 1990), Hansen and Sargent (2005) interpret this recursive implementation of the commitment problem as one in which as time unfolds the decision maker's

[9] In Hansen and Sargent (2005) there are typos in Eq. (50) ($\tilde{\Delta}_{k-1}$ should be included as an argument of the function M on the right) and in the expression for $\hat{z}_t$ in Section 7.5 (Δ_t should be $\check{\Delta}_t$) that correspond to Eqs. (40) and (38), respectively. Appendix B describes an alternative formulation of these recursions.

[10] The distortion associated with $\check{m}_j$ implies a step in updating beliefs that is in addition to the updating associated with the ordinary Kalman filter defined in (9), (10), and (11) to update the hidden state conditional mean of the hidden state. Since $\check{m}_j$ is an exponential of a quadratic function of z_j, these distortions are computed using the normal density and a complete the square argument.

benchmark model changes in ways that depend on actions that affected past values of the one-period objective function. That reflects the feature that the Kalman filtering Eqs. (10)–(11) are backward-looking.

The wish to acquire robust estimators leads one to explore the utility consequences of distorting the evolution of hidden states. Under commitment, the date zero utility function is the relevant one for inducing robustness via exponential twisting of probabilities of hidden states. The change in benchmark models represented in steps (38) and (41) captures this.

As Whittle (1990) emphasized, the decision rule (39) has forward-looking components that come from 'control' and backward-looking components that come from 'filtering under commitment'. The sufficient statistic $\check{z}_t$ used as a benchmark in state estimation is backward-looking. When hidden state variables enter the one-period utility function, $\check{z}_t$ can deviate from the state estimate obtained by direct application of the Kalman filter. The forward-looking component comes from the control component of the problem through the matrices Ω_{t+1} and Ω_t in (37). We combine both components to express a robust action partly as a function of a distorted estimate of the hidden state $\hat{z}_t$.

Example 8.1 (*Hidden state not in objective*). Suppose that P_2=0, R_{12}=0, and R_{22}=0, so that the one-period objective does not depend on the hidden state. In this case, there is an alternative way to solve the robust control problem that first solves the filtering problem and then computes an ordinary robust control for the reduced information configuration associated with the innovations representation (9)–(11).[11]

Write the solution to the ordinary (non-robust) filtering problem as

$$\overline{z}_{t+1} = A_{21}y_t + A_{22}\overline{z}_t + B_2 a_t + K_2(\Delta_t)\overline{w}_{t+1},$$

where the innovation

$$\overline{w}_{t+1} = D_2(z_t - \overline{z}_t) + G w_{t+1}$$

is normal with mean zero and covariance matrix

$$D\Sigma(\Delta)_t D' + GG'$$

conditioned on $\mathcal{S}_t$. Instead of distorting the joint distribution (w_{t+1}, x_t), we can distort the distribution of the innovation $\overline{w}_{t+1}$ conditioned on $\mathcal{S}_t$. It suffices to add a distortion $\overline{v}_t$ to the mean of $\overline{w}_{t+1}$ with entropy penalty

$$\theta \overline{v}_t'(D\Sigma(\Delta_t)D' + GG')^{-1}\overline{v}_t,$$

and where $\overline{v}_t$ is restricted to be a function of the signal history. While the conditional covariance is also distorted, certainty equivalence allows us to compute the mean distortion by solving a deterministic zero-sum two-player game. As in the robustness problem with full information, discounting causes concerns about robustness to wear off over time.

Remark 8.2. Under the Example 8.1 assumptions, the only difference between Games III and IV is that Game III discounts time t contributions to entropy while Game IV does not. When the objective function satisfies the special conditions of Example 8.1 and $\beta = 1$, outcomes of Games III and IV coincide.

Example 8.3 (*Pure estimation*). The state is exogenous and unaffected by the control. The objective is to estimate $-Px_t$. The control is an estimate of $-Px_t$. To implement this specification, we set B=0, Q=I, and $R=P'P$. For this problem, the solution of (37) for Game IV is $\Omega_t = 0$ for all $t \geq 0$ because $a = -Px$ sets the full information objective to zero. The solution to the estimation problem is $a_t = -P\check{x}_t$ where $\check{x}_t = [\begin{smallmatrix} y_t \\ \check{z}_t \end{smallmatrix}]$ and $\tilde{z}_t = \check{z}_t = \hat{z}_t$. In this case, the Game IV recursions (38)–(43) collapse to

$$\check{z}_t = \mathsf{M}(y_{t-1}, \check{z}_{t-1}, a_{t-1}, s_t, D(\check{\Delta}_{t-1})), \tag{45}$$

$$\check{\Delta}_t = \mathsf{C} \circ \mathsf{D}(\check{\Delta}_{t-1}), \tag{46}$$

where $\mathsf{D}(\check{\Delta}_t) = [(\check{\Delta}_t)^{-1} - (\beta^k/\theta)R_{22}]^{-1}$ is the operator affiliated with (40). For $\beta = 1$, it can be verified that these are the recursions described by Hansen and Sargent (2008, Chapter 17).

9. Calibrating (θ_1, θ_2) with detection error probabilities

Following Pratt (1964), Ljungqvist and Sargent (2004, pp. 424–426) describe how to calibrate a coefficient of relative risk aversion via a mental experiment that asks a person how much he or she would be willing to pay to avoid a gamble with known probability. But risk with a known distribution differs from uncertainty about the prevailing distribution. Our multiplier parameters θ_1 and θ_2 express a decision maker's distress at not knowing probabilities, in particular, the probability distribution of next period's state z^* given z in the case of θ_1, and the probability distribution of this period's state z, in the case of θ_2. In settings with an observed state vector, Anderson et al. (2003) and Hansen and Sargent (2008, Chapter 9) have described a method for calibrating θ_1 by using a mental experiment cast in terms of Bayesian detection

[11] This approach is used in the asset pricing applications in Hansen and Sargent (2010).

error probabilities. In this section, we describe how to extend that general approach in a way that allows us to calibrate θ_2 as well.

An equilibrium of one of our two-player zero-sum games can be represented in terms of the following law of motion for $(y,\check{z},s)$:

$$y^* = A_{11}y + A_{12}\check{z} + B_1 a + C_1 w^* + A_{12}(z-\check{z}),$$

$$\check{z}^* = A_{21}y + A_{22}\check{z} + B_2 a + K_2(\Delta)Gw^* + K_2(\Delta)D_2(z-\check{z}),$$

$$s^* = D_1 y + D_2\check{z} + Ha + Gw^* + D_2(z-\check{z}),$$

$$\Delta^* = \mathcal{C}(\Delta), \tag{47}$$

where under the approximating model

$$w^* \sim \mathcal{N}(0,I) \quad \text{and} \quad z-\check{z} \sim \mathcal{N}(0,\Delta), \tag{48}$$

and under the worst-case model associated with a (θ_1,θ_2) pair

$$w^* \sim \mathcal{N}(\tilde{v},\Sigma(\Delta)) \quad \text{and} \quad z-\check{z} \sim \mathcal{N}(u,\Gamma(\Delta)). \tag{49}$$

We have shown how to compute decision rules for a, $\tilde{v}$, u, $\Sigma(\Delta)$, $\Gamma(\Delta)$ for each our zero-sum two-player games; $a,\tilde{v},u$ are linear functions of $y,\check{z}$. Evidently, under the approximating model

$$s^* \sim \mathcal{N}(\overline{D}_1 y + \overline{D}_2\check{z},\Omega_a), \tag{50}$$

where $\Omega_a(\Delta) = GG' + D_2\Delta D_2'$ and the $(\bar{\cdot})$ over a matrix indicates that the feedback rule for a has been absorbed into that matrix; while under the worst-case model

$$s^* \sim \mathcal{N}(\hat{D}_1 y + \hat{D}_2\check{z},\Omega_w), \tag{51}$$

where $\Omega_w(\Delta) = G\Sigma(\Delta)G' + D_2\Gamma(\Delta)D_2'$ and the $(\hat{\cdot})$ over a matrix indicates the feedback rules for a and the conditional means $\tilde{v},u$ have been absorbed into that matrix.

Where N is the number of variables in s_{t+1}, conditional on $y_0,\check{z}_0$, the log likelihood of $\{s_{t+1}\}_{t=0}^{T-1}$ under the approximating model is

$$\log L_a = -\frac{1}{T}\sum_{t=0}^{T}\left[\frac{N}{2}\log(2\pi) + \frac{1}{2}\log|\Omega_a(\Delta_t)| + \frac{1}{2}(s_{t+1}-\overline{D}_1 y_t - \overline{D}_2\check{z}_t)'\Omega_a(\Delta_t)^{-1}(s_{t+1}-\overline{D}_1 y_t - \overline{D}_2\check{z}_t)\right] \tag{52}$$

and the log likelihood under the worst-case model is

$$\log L_w = -\frac{1}{T}\sum_{t=0}^{T}\left[\frac{N}{2}\log(2\pi) + \frac{1}{2}\log|\Omega_w(\Delta_t)| + \frac{1}{2}(s_{t+1}-\hat{D}_1 y_t - \hat{D}_2\check{z}_t)'\Omega_w(\Delta_t)^{-1}(s_{t+1}-\hat{D}_1 y_t - \hat{D}_2\check{z}_t)\right]. \tag{53}$$

By applying procedures like those described in Hansen et al. (2002) and Anderson et al. (2003), we can use simulations in the following ways to approximate a detection error probability:

- Repeatedly simulate $\{y_{t+1},\check{z}_{t+1},s_{t+1}\}_{t=0}^{T-1}$ under the approximating model. Use (52) and (53) to evaluate the log likelihood functions of the approximating model and worst case model for a given (θ_1,θ_2). Compute the fraction of simulations for which $\log L_w > \log L_a$ and call it r_a. This approximates the probability that the likelihood ratio says that the worst-case model generated the data when the approximating model actually generated the data.
- Repeatedly simulate $\{y_{t+1},\check{z}_{t+1},s_{t+1}\}_{t=0}^{T-1}$ under the worst-case model affiliated with a given (θ_1,θ_2) pair. Use (52) and (53) to evaluate the log likelihood functions of the approximating and worst case models. Compute the fraction of simulations for which $\log L_a > \log L_w$ and call it r_w. This approximates the probability that the likelihood ratio says that the approximating model generated the data when the worst-case model generated the data.
- As in Hansen et al. (2002) and Anderson et al. (2003), define the overall detection error probability to be

$$p(\theta_1,\theta_2) = \frac{1}{2}(r_a + r_w). \tag{54}$$

9.1. Practical details

The detection error probability $p(\theta_1,\theta_2)$ in (54) can be used to calibrate the pair (θ_1,θ_2) jointly. This seems to be the appropriate procedure for Game II, especially when z does not appear in the objective function. However, for Game I, we think that the following sequential procedure makes sense.

1. First pretend that y, z are both observable. Calibrate θ_1 by calculating detection error probabilities for a system with an observed state vector using the approach of Hansen et al. (2002) and Hansen and Sargent (2008, Chapter 9).
2. Then having pinned down θ_1 in step 1, use the approach leading to formula (54) to calibrate θ_2.

This procedure takes the point of view that θ_1 measures how difficult it would be to distinguish one model of the partially hidden state from another if we were able to observe the hidden state, while θ_2 measures how difficult it is to distinguish alternative models of the hidden state. The probability $p(\theta_1,\theta_2)$ measures both sources of model uncertainty.

10. Concluding remarks

Ellison and Sargent (2009) use a version of our Game I to interpret a puzzle spotted by Romer and Romer (2008), who found that although FOMC members know Federal Reserve Board staff forecasts of inflation and unemployment before making their own distinct forecasts: (1) the FOMC makes systematically worse forecasts than the staff that are biased in particular directions and (2) the discrepancies between the FOMC's forecasts and the staff's influence FOMC decisions. Romer and Romer deplore finding 2 and suggest that the FOMC should leave forecasting to the staff and not allow the gap between FOMC forecasts and the staff's to influence FOMC decisions. Ellison and Sargent interpret Romer and Romer's judgement that FOMC members should leave forecasting to the staff as a recommendation to trust the Board staff's model. Using a calibrated model that captures some tradeoffs that preoccupy the FOMC and that features hidden state variables that seem to influence the FOMC, including in particular a 'natural rate of unemployment', Ellison and Sargent reinterpret Romer and Romer's empirical findings by assuming that the FOMC fears that the staff's model is misspecified and wants decisions that are robust to misspecifications. Ellison and Sargent find that to explain Romer and Romer's findings, it is very helpful to assume that the FOMC is most concerned about misspecification of its statistical model of the natural rate of unemployment. This is why they use a version of Game I that features the natural rate of unemployment as a hidden state variable that the FOMC cares about vitally.

Hansen and Sargent (2010) extend a version of Game I to study asset pricing in a long-run risk model by positing two submodels over which a representative consumer puts a prior distribution that he does not fully trust. In each submodel, the hidden state vector pins down the conditional mean rate of consumption growth as the sum of an unknown constant and a first-order autoregressive process. For one of the submodels, the first-order autoregressive process has little serial correlation, while for the other, called the long-run risk model, that process is highly persistent. Two such models are known to be difficult to distinguish for U.S. consumption growth data, even with fairly long time series. For each submodel, the representative consumer estimates the hidden state via a Kalman filter, but then perturbs its distribution to adjust for fear of misspecification, as described in this paper. Hansen and Sargent (2010) also show how to adjust the Bayesian posterior over the two submodels for distrust. They study how the resulting adjustments for fear of misspecification impinge on what are usually interpreted as market prices of risk, but which they instead interpret as market prices of model uncertainty.

The Hansen and Sargent (2010) model illustrates how we can step slightly outside the LQG structure of this paper to consider more general settings in which an additional hidden Markov state indexes a finite set of LQG submodels, for example, as in the models without fear of model misspecification analyzed by Svensson and Williams (2008).[12] It is possible to use the ideas and computations in this paper to adapt the Svensson and Williams structure to incorporate fears of misspecification of the submodels and of the distribution over submodels. That would involve calculations closely related to ones that Hansen and Sargent (2010) use to model countercyclical uncertainty premia and that Cogley et al. (2008) use to design a robust monetary policy when part of the monetary policy design problem is to affect the evolution of submodel probabilities through purposeful experimentation.

Acknowledgments

We thank the editor, two referees, Francisco Barillas, Saki Bigio, Martin Ellison, Christian Matthes, and Leandro Nascimento for helpful comments. Hansen and Sargent thank the National Science Foundation for research support. Sargent thanks the Bank of England for providing research time and support while he was a Houblon–Norman fellow.

Appendix A. Modified certainty equivalence

Connections between the outcomes of the following two problems allow us to compute the T^1 operator easily.

Problem Appendix A.1. Let $V(x) = -\frac{1}{2}x'\Omega x - \omega$, where Ω is a positive definite symmetric matrix. Consider the control problem

$$\min_{v} V(x^*) + \frac{\theta}{2}|v|^2$$

subject to a linear transition function $x^* = Ax + Cv$. If θ is large enough that $I - \theta^{-1}C'\Omega C$ is positive definite, the problem is well posed and has solution

$$v = Kx, \tag{A.1}$$

$$K = [\theta I - C'\Omega C]^{-1}C'\Omega A. \tag{A.2}$$

[12] The periodic models of Hansen and Sargent (2007a, Chapter 17) are closely related to the structures of Svensson and Williams (2008).

The following problem uses (A.1) and (A.2) to compute the T^1 operator:

Problem Appendix A.2. For the same value function $V(x) = -\frac{1}{2}x'\Omega x - \omega$ that appears in Problem Appendix A.1, let the transition law be

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

where $w^* \sim \mathcal{N}(0,I)$. The T^1 operator gives the indirect utility function of the following minimization problem:

$$\min_{m^*} E[m^* V(x^*) + \theta m^* \log m^*].$$

The minimizer is

$$m^* \propto \exp\left(\frac{-V(x^*)}{\theta}\right) = \exp\left[-\frac{1}{2}(w^*-\nu)'\Sigma^{-1}(w^*-\nu) + \frac{1}{2}w^* \cdot w^* - \frac{1}{2}\log\det\Sigma\right],$$

where ν is given by (A.1) and (A.2) from Problem Appendix A.1, the worst-case variance $\Sigma = (I - \theta^{-1}C'\Omega C)^{-1}$, and the entropy of m^* is

$$Em^* \log m^* = \tfrac{1}{2}[|\nu|^2 + \text{trace}(\Sigma - I) - \log\det\Sigma].$$

Therefore, we can compute the objects (ν, Σ) needed to form T^1 by solving the deterministic Problem Appendix A.1.

Appendix B. Alternative formulation

When Δ_t is non-singular, recursions (40)–(43) can be implemented with the following recursions that are equivalent to the formulation of Başar and Bernhard (1995). Let

$$\check{A} \doteq A_{22} - C_2 G'(GG')^{-1} D_2,$$

$$\check{N} \doteq C_2 C_2' - C_2 G'(GG')^{-1} G C_2'.$$

Then we can attain $\check{z}_t$ directly via the recursions:

$$\check{\Delta}_{t+1} = \check{A}\left[(\check{\Delta}_t)^{-1} - \frac{\beta^t}{\theta}R_{22} + D_2'(GG')^{-1}D_2\right]\check{A}' + \check{N} \tag{B.1}$$

and

$$\check{z}_{t+1} = A_{21}y_t + \check{A}\check{z}_t + B_2 a_t + \check{A}\left[(\check{\Delta}_t)^{-1} - \frac{\beta^t}{\theta}R_{22} + D_2'(GG')^{-1}D_2\right]^{-1} D_2'(GG')^{-1}(s_{t+1} - D_1 y_t - D_2\check{z}_t)$$
$$+ \frac{\beta^t}{\theta}\check{A}\left[(\check{\Delta}_t)^{-1} - \frac{\beta^t}{\theta}R_{22} + D_2'(GG')^{-1}D_2\right]^{-1}(P_2' a_t + R_{12}' y_t + R_{22}\check{z}_t). \tag{B.2}$$

Appendix C. Matlab programs

This appendix describes how to use Matlab programs that solve and simulate outcomes of our four games.

Four object oriented programs named

```
PreData_ComputeGameOne.m
PreData_ComputeGameTwo.m
PreData_ComputeGameThree.m
PreData_ComputeGameFour.m
```

compute objects in the respective games that do not depend on data. After these objects have been computed, the programs

```
ComputeGameOne_dataparts.m
ComputeGameTwo_dataparts.m
ComputeGameThree_dataparts.m
ComputeGameFour_dataparts.m
```

generate time series of decisions and filtered estimates.

A sample driver files illustrates two different ways of using these Game solving functions with the context of a permanent income example. The program `example_permanent_income.m` simulates data from the approximating model and n solves Game I. To illustrate how to use an external data set, it saves and loads the simulated data and proceeds to solve each of Games II–IV. Of course it is also possible to simulate a new data set each time you call a game solving function. Examples of this last feature are included as comments in the driver file. Once the functions mentioned above have computed the objects of interest, you can extract them for further analysis. Both the accompanying readme file and the permanent income driver file show you ways to extract these results. Finally, we have provided a couple of files that plot subsets of these results.

References

Anderson, E., Hansen, L., Sargent, T., 2003. A quartet of semigroups for model specification, robustness, prices of risk, and model detection. Journal of the European Economic Association 1 (1), 68–123.

Başar, T., Bernhard, P., 1995. Optimal Control and Related Minimax Design Problems H^{∞}, second ed. Birkhauser.

Cappé, O., Moulines, R., Rydén, T., 2005. Inference in Hidden Markov Models. Springer.

Cerreia, S., Maccheroni, F., Marinacci, M., Montrucchio, L., 2008. Uncertainty averse preferences. Working Paper, Collegio Carlo Alberto.

Cogley, T., Colacito, R., Hansen, L.P., Sargent, T.J., 2008. Robustness and U.S. monetary policy experimentation. Journal of Money, Credit and Banking 40 (8), 1599–1623.

Elliott, R.J., Aggoun, L., Moore, J.B., 1995. Hidden Markov Models: Estimation and Control. Springer-Verlag, New York.

Ellison, M., Sargent, T.J., 2009. A defence of the FOMC. Economics series Working Papers, University of Oxford, Department of Economics.

Epstein, L., Schneider, M., 2003. Recursive multiple priors. Journal of Economic Theory 113 (1), 1–31.

Epstein, L.G., Schneider, M., 2007. Learning under ambiguity. Review of Economic Studies 74 (4), 1275–1303.

Hamilton, J.D., 1989. A new approach to the economic analysis of nonstationary time series and the business cycle. Econometrica 57 (2), 357–384.

Hansen, L.P., Sargent, T.J., 2001. Robust control and model uncertainty. American Economic Review 91 (2), 60–66.

Hansen, L.P., Sargent, T.J., 2005. Robust estimation and control under commitment. Journal of Economic Theory 124 (2), 258–301.

Hansen, L.P., Sargent, T.J., 2007a. Recursive models of dynamic linear economies. Unpublished monograph, University of Chicago and Hoover Institution.

Hansen, L.P., Sargent, T.J., 2007b. Recursive robust estimation and control without commitment. Journal of Economic Theory 136, 1–27.

Hansen, L.P., Sargent, T.J., 2008. Robustness. Princeton University Press, Princeton, NJ.

Hansen, L.P., Sargent, T.J., 2010. Fragile beliefs and the price of uncertainty. Quantitative Economics, in press.

Hansen, L.P., Sargent, T.J., Wang, N.E., 2002. Robust permanent income and pricing with filtering. Macroeconomic Dynamics 6, 40–84.

Kasa, K., 2001. A robust Hansen–Sargent prediction formula. Economics Letters 71 (1), 43–48.

Kasa, K., 2002. Model uncertainty, robust policies, and the value of commitment. Macroeconomic Dynamics 6 (01), 145–166.

Ljungqvist, L., Sargent, T.J., 2004. Recursive Macroeconomic Theory, second ed. MIT Press, Cambridge, MA.

Pratt, J.W., 1964. Risk aversion in the small and in the large. Econometrica 32, 122–136.

Romer, C.D., Romer, D.H., 2008. The FOMC versus the staff: where can monetary policymakers add value? American Economic Review 98 (2), 230–235

Sclove, S.L., 1983. Time-series segmentation: a model and a method. Information Sciences 29 (1), 7–25.

Sims, C.A., Zha, T., 2006. Were there regime switches in U.S. monetary policy? American Economic Review 96 (1), 54–81

Svensson, L.E., Williams, N., 2008. Optimal monetary policy under uncertainty in dsge models: a markov jump-linear-quadratic approach. NBER Working Papers, National Bureau of Economic Research, Inc.

Whittle, P., 1990. Risk-Sensitive Optimal Control. John Wiley & Sons, New York.