Recursive Formulations of Robust Estimation and Control Without Commitment

Lars Peter Hansen

Thomas J. Sargent

University of Chicago and NYU

Stochastic Versions of Robust Control: Examples

- Petersen, James, and Dupuis (2000) *IEEE Transactions in Automatic Control*
- Hansen, Sargent, Turmuhambetova, and Williams (2004) forthcoming in *JET*

Related work by Epstein and Schneider (2003a) and Epstein and Schneider (2003b)

Motivation for research

- Alternative time series models are hard to distinguish statistically, but can have important differences for valuation and assessment of the impact of uncertainty
- Probability models typically used are arguably approximations where approximation errors are challenging to pose in a full probabilistic manner

Current paper:

- Study robust notions of learning and model averaging in conjunction with decision making.
- *Recursive* formulation - related paper considers a *commitment* formulation - Hansen-Sargent forthcoming in *JET*

Exploit well known algorithms for updating probabilities

Hidden state Markov processes

- Kalman filtering - Linear, Gaussian models
- Wonham filtering - Finite state hidden state Markov chain disguised by Brownian motion
- Zackai equation - clever change in probabilities
- particle filtering

Learning and Permanent Income Model

- LQ permanent income model with a hidden states.
- One of the states is an unknown time invariant parameter - mean of labor income
- Application of Kalman filtering.
- explore sensitivity of decisions to posterior distributions

Learning about consumption growth rate: Kalman filter

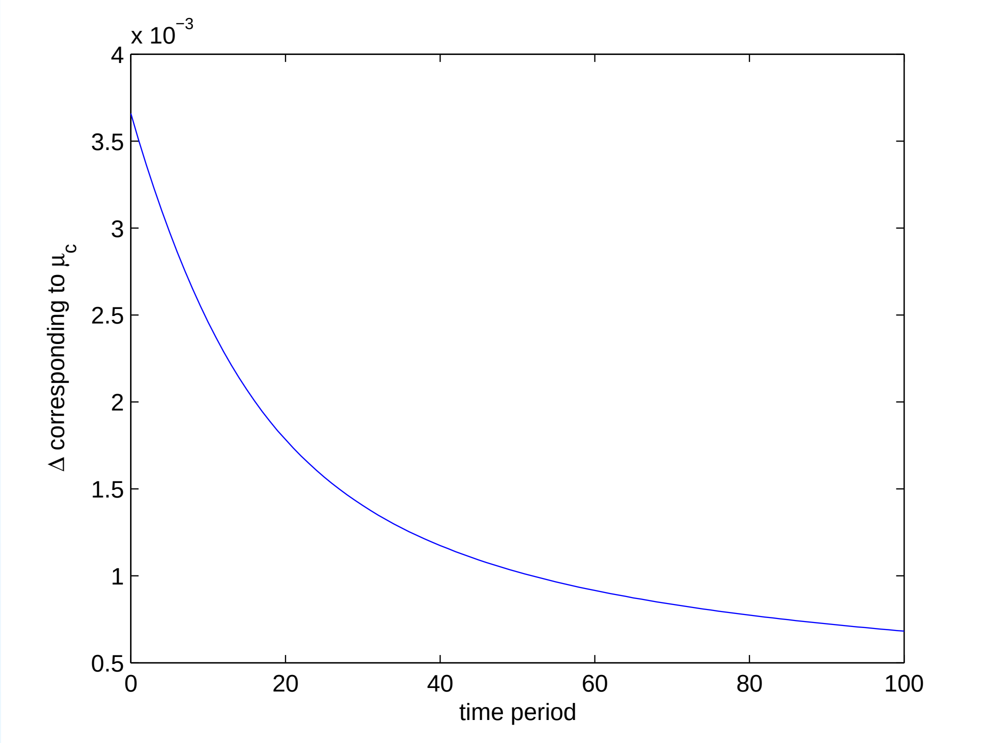

Time series computed posterior volatilities of unknown growth rates.

Conditional volatility of consumption

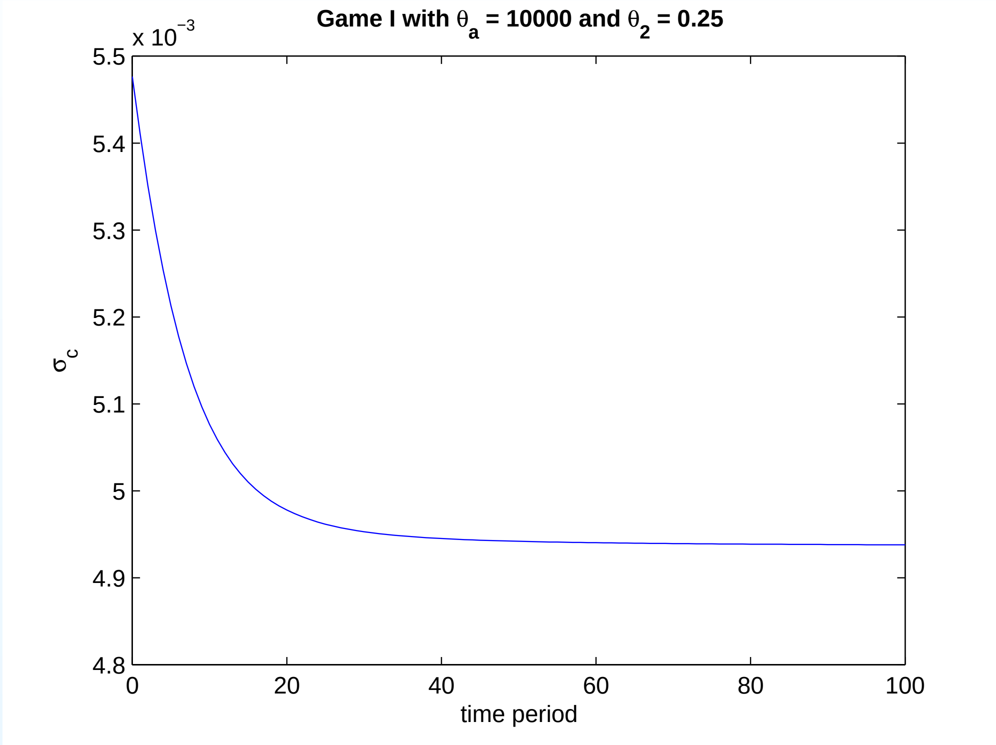

Posterior volatility for consumption

Relative consumption distortion

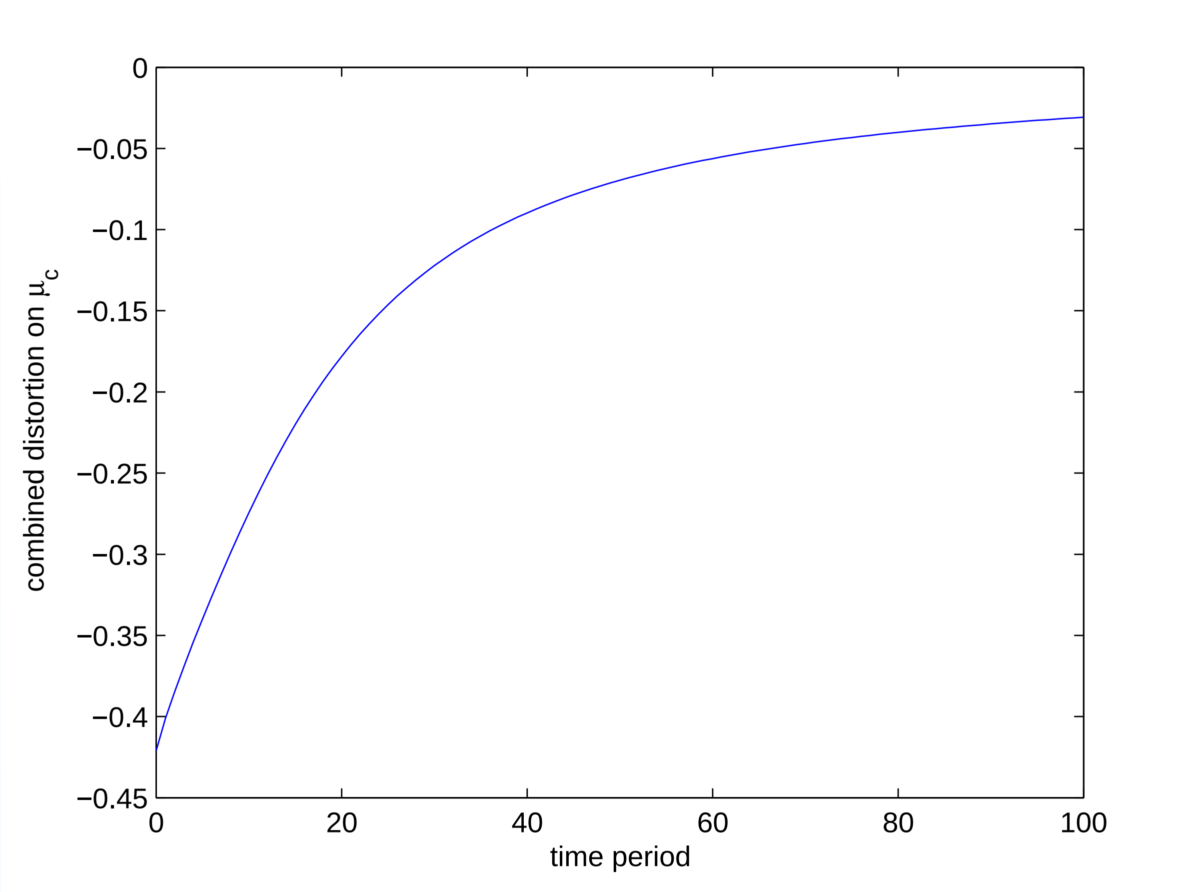

Mean distortion in consumption relative to the volatility

Stochastic growth model with hidden growth states

- Technology shock process in which growth states change according to a Markov process.
- Wonham filter to produce posterior probabilities.
- Posterior sensitivity analysis

Probability of High Growth State: Wonham filter

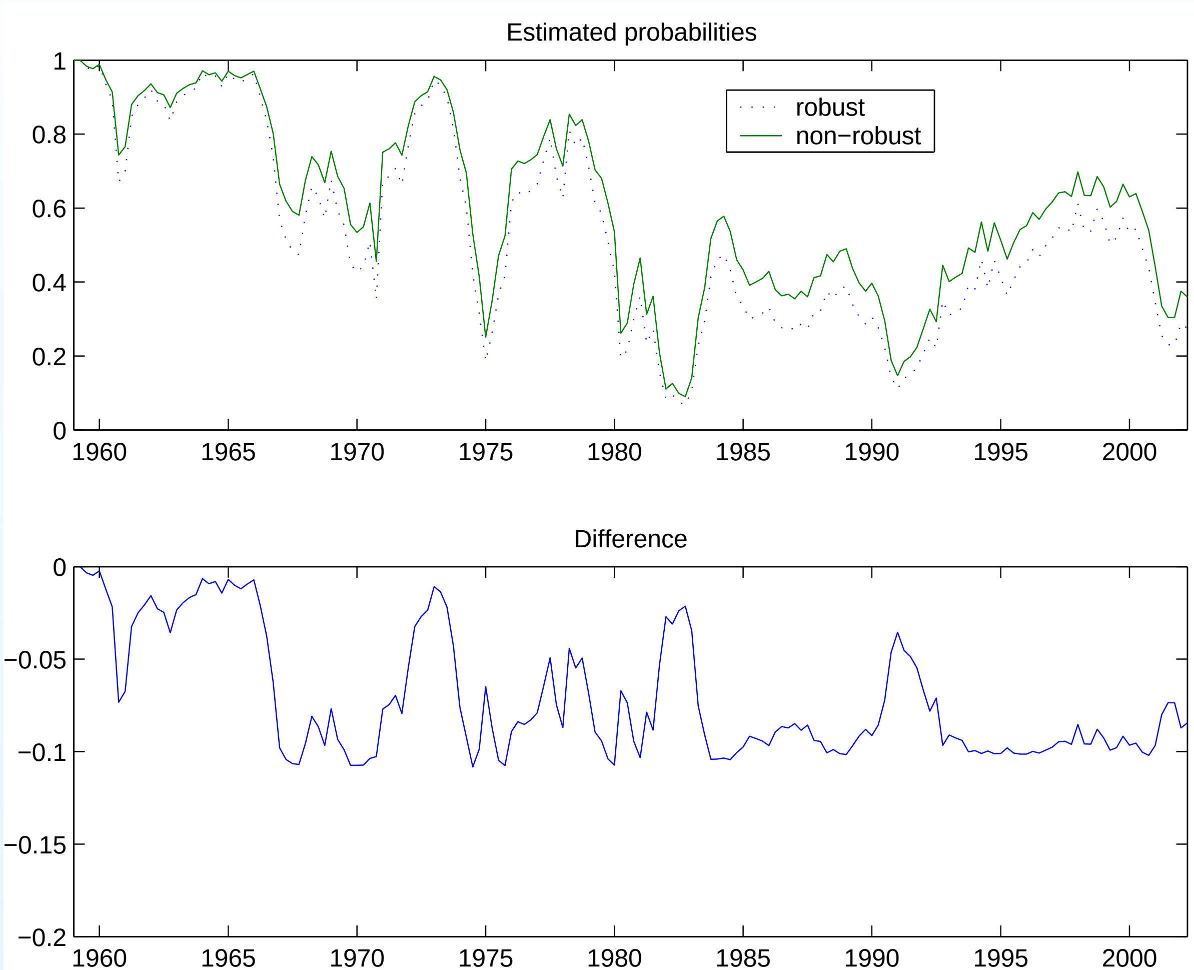

Top graph: probability of the high growth state

Bottom graph: difference between the two series.

Uncertainty premia in hidden growth model

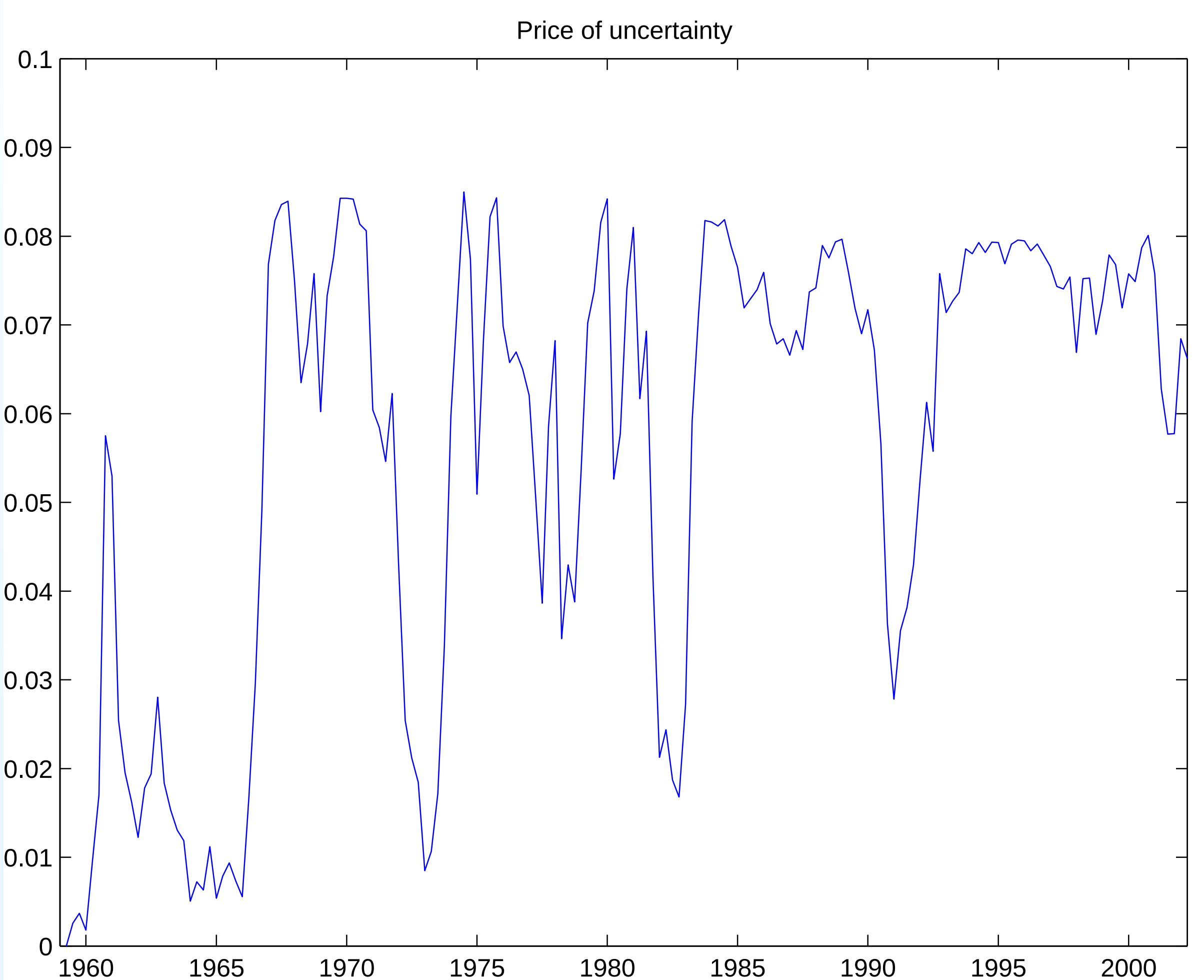

Implied distortions in the consumption distribution

Model selection and long run risk

- Variety of recent work in macro/asset pricing studies implicitly implications of models with similar statistical properties but different long run properties.
- Prices in decentralized economies can be sensitive to these changes.
- Why focus on one of the competing models?
- Robust model selection/averaging

Martingales and Distorted Probabilities

- Event collections $\{\mathcal{X}_t : t \geq 0\}$ where $\mathcal{X}_t$ is the date t information set. Let Pr denote a probability measure on $\mathcal{X}_\infty = \vee_{t\geq 0}\mathcal{X}_t$.
- Nonnegative martingale $\{M_t : t \geq 0\}$ where $M_0 = 1$. In particular, $E\left(M_t|\mathcal{X}_0\right) = 1$.
- Distorted probability

$$\tilde{E}(x_t|\mathcal{X}_0) = E\left(M_t x_t|\mathcal{X}_0\right)$$

where M_t is a likelihood ratio or a R-N derivative.

Entropy

- What is it?

  $$EM \log M$$

  where $EM = 1$. Average log - likelihood.
  Use gradient inequality

  $$M \log M \geq M - 1$$

- Why the name?

  > *When Shannon had invented his quantity and consulted von Neumann on what to call it, von Neumann replied: ‘Call it entropy. It is already in use under that name and besides, it will give you a great edge in debates because nobody knows what entropy is anyway.’*

Entropy penalization problem

- Problem

$$\min_{M \geq 0, EM=1} E\left(M[V + \theta \log(M)]\right)$$

- Solution - exponential tilting

$$M^* = \frac{\exp\left(-\frac{1}{\theta}V\right)}{E\left[\exp\left(-\frac{1}{\theta}V\right)\right]}$$

- Minimized objective

$$-\theta \log E\left[\exp\left(-\frac{1}{\theta}V\right)\right]$$

Static formulation of robust control

- Let a in an action in a feasible set $\mathcal{A}$ and let x an unknown state. Informational restrictions may be imposed on the action. Actions and states can be processes. Objective can be discounted utility.
- Problem

$$\sup_{a \in \mathcal{A}} \inf_{M \geq 0, EM=1} E\left(M[V(a,x) + \theta \log(M)]\right)$$

Worst case M^* depends on action. Zero sum game.

- Reverse orders:

$$\inf_{M \geq 0, EM=1} \sup_{a \in \mathcal{A}} E\left(M[V(a,x) + \theta \log(M)]\right)$$

Action a^* optimizes against a fixed probability. 'Bayesian solution'.

Multiplicative decomposition

Form m_{t+1}:

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

where

$$E\left(m_{t+1} | \mathcal{X}_t\right) = 1$$

Then

$$M_t = \prod_{j=1}^{t} m_j$$

The random variable m_{t+1} distorts the transition density between date t and date $t+1$. Factoring (relative) likelihood.

Robust Control and Discounted Entropy

- Recursive solution in which date t minimizing agent chooses m_{t+1} subject to penalty $\theta E[m_{t+1} \log(m_{t+1})]$.
- Hansen and Sargent (1995),
  Anderson, Hansen, and Sargent (2003) and
  Hansen, Sargent, Turmuhambetova, and Williams (2004)

Complaints

- Single benchmark model with perturbations around that model - Levin, Wieland and Williams
- No focused consideration of parameter uncertainty - Onatski and Williams
- No scope for learning

Introduce a hidden state Markov process as a motivation for learning.

Basic Idea

Two robustness recursions:

1. Allow for misspecified dynamics as before using m_{t+1} conditioned on a big information set that includes information on a history of hidden states.
2. Sensitivity analysis of posterior probabilities used for averaging over the hidden states.

Observations

1. hidden states can be time invariant and hence index alternative models or unknown parameters
2. hidden states can evolve as a Markov chain - time varying parameter models - regime shift models
3. exploit tools for solving hidden state Markov chain models
4. minimizing agent has an informational advantage
5. recursive

Learning, commitment and recursivity

Related papers:

- Hansen and Sargent (2005) commitment counterpart to this paper - builds on Basar and Bernhard (1995) and others.
- Chamberlain (2000) and Knox (2003) closely related commitment problems.
- Epstein and Schneider (2003a) and
  Epstein and Schneider (2003b) avoid commitment in a different but related formulation

Alternative decompositions of M_t

Let $\mathcal{S}_t$ denote a signal history smaller that $\mathcal{X}_t$.

- Decompose $M_{t+1} = m_{t+1} M_t$. Use m_{t+1} to distort dynamics conditioned on the hidden state history.
- Use entropy penalty $E[m_{t+1} \log(m_{t+1}) | \mathcal{X}_t] = 1$.
- Decompose:

$$M_t = h_t G_t, E(h_t | \mathcal{S}_t) = 1.$$

- Use h_t to distort the probabilities assigned to $\mathcal{X}_t$ events conditioned on $\mathcal{S}_t$
- Use an entropy penalty on h_t:

$$E(h_t \log h_t | \mathcal{S}_t)$$

Baseline Problem Setup

- Partition a state vector as $x_t = \begin{bmatrix} y_t \\ z_t \end{bmatrix}$, where y_t is observed and z_t is not. Let s_t denote a vector of signals of the unobserved state z_t.
- Let Z denote a space of admissible unobserved states, $\mathcal{Z}$ a corresponding sigma algebra of subsets of states, and λ a measure on the measurable space of hidden states $(Z, \mathcal{Z})$. Let S denote the space of signals, $\mathcal{S}$ a corresponding sigma algebra, and η a measure on the measurable space $(S, \mathcal{S})$ of signals.

State Evolution

- Signals and states are determined by the transition functions

$$
\begin{aligned}
y_{t+1} &= \pi_y(s_{t+1}, y_t, a_t), & (1)\\
z_{t+1} &= \pi_z(x_t, a_t, w_{t+1}), & (2)\\
s_{t+1} &= \pi_s(x_t, a_t, w_{t+1}) & (3)
\end{aligned}
$$

  where $\{w_{t+1} : t \geq 0\}$ is an i.i.d. sequence of random vectors.
- Observable state evolution:

$$y_{t+1} = \bar{\pi}_y(x_t, a_t, w_{t+1}).$$

- Equations (3) and (2) determine a conditional density $\tau(z_{t+1}, s_{t+1} | x_t, a_t)$ relative to the product measure $\lambda \times \eta$.

Recursive Formulation

- Use τ to construct two densities for the signal:

$$\begin{aligned} \kappa(s^*|y_t, z_t, a_t) &\doteq \int \tau(z^*, s^*|y_t, z_t, a_t) d\lambda(z^*) \\ \varsigma(s^*|y_t, q_t, a_t) &\doteq \int \kappa(s^*|y_t, z, a_t) q_t(z) d\lambda(z). \end{aligned}$$

- By Bayes' rule,

$$q_{t+1}(z^*) = \frac{\int \tau(z^*, s_{t+1}|y_t, z, a_t) q_t(z) d\lambda(z)}{\varsigma(s_{t+1}|y_t, q_t, a_t)} \doteq \pi_q(s_{t+1}, y_t, q_t, a_t).$$

π_q can be computed by using filtering methods that specialize Bayes' rule (*e.g.*, the Kalman filter or a discrete time version of the Wonham filter or Zackai equation.

Constructed State Vector

- Take (y_t, q_t) as the state with transition law

$$
\begin{aligned}
y_{t+1} &= \pi_y(s_{t+1}, y_t, a_t) \\
q_{t+1} &= \pi_q(s_{t+1}, y_t, q_t, a_t).
\end{aligned}
$$

- Choose a_t as a function of (y_t, q_t).

One strategy is apply our earlier ”full information” approach to stochastic robust control to this problem!! We will consider alternatives.

Robustness and hidden states

- Two agents face different information restrictions. Minimizing agent can find distortions conditioned on hidden states.
- Break link between recursive and commitment formulations using two robustness recursions.
- Control is forward looking and solved by backward induction. Prediction is backward looking and solved by forward induction. Tension in the construction of worst-case models.

Recursive formulation

- Two distinct distortions at date t:
  1. Distort dynamics for (x_{t+1}, s_{t+1}) conditioned on (x_t, q_t). m_{t+1} distortion from before. Distort probabilities assigned to $\mathcal{X}_{t+1}$ conditioned on $\mathcal{X}_t$.
  2. Distort hidden state probabilities q_t or more generally the conditional probabilities assigned to $\mathcal{X}_t$ events conditioned on $\mathcal{S}_t$.
- We do not simply ”reduce compound lotteries” where the compounding is over hidden state z_t.

Recursive Distortion of State Dynamics

Consider a value function $V(y_{t+1}, q_{t+1}, z_{t+1})$,

$$\mathsf{T}^1(V|\theta)(y,q,z,a) =$$
$$-\theta \log \int \exp\left(-\frac{V[\pi(s^*, y, q, a), z^*]}{\theta}\right) \tau(z^*, s^*|y, z, a) d\lambda(z^*) d\eta(s^*).$$

The transformation T^1 maps a value function that depends on the state (y, q, z) into a risk-adjusted value function that depends on (y, q, z, a). Associated with this *risk adjustment* is a worst-case distortion in the transition dynamics for the state and signal process:

$$\phi_t(z^*, s^*) = \frac{\exp\left(-\frac{V[\pi(s^*, y_t, q_t, a_t), z^*]}{\theta}\right)}{E\left[\exp\left(-\frac{V[\pi(s_{t+1}, y_t, q_t, a_t), z_{t+1}]}{\theta}\right) |\mathcal{X}_t\right]}.$$

Recursive Distortion of State Probabilities

Consider a value function of the form: $\hat{V}(y_t, q_t, z_t, a_t)$ and operator:

$$\mathsf{T}^2(\hat{V}|\theta)(y,q,a) = -\theta \log \int \exp\left[-\frac{\hat{V}(y,q,z,a)}{\theta}\right] q(z) d\lambda(z).$$

The worst case density conditioned on $\mathcal{S}_t$ is $\psi_t(z) q_t(z)$ where

$$\psi_t(z) = \frac{\exp\left(-\frac{\hat{V}(y_t,q_t,z,a_t)}{\theta}\right)}{E\left[\exp\left(-\frac{\hat{V}(y_t,q_t,z,a_t)}{\theta}\right)|\mathcal{S}_t\right]}.$$

Recursive Game I

Consider an approach that keeps track of a value function that depends on the hidden state.

$$\check{W}(y, q, z) = U(x, a) + \mathsf{T}^1 \left[\beta \check{W}^*(y^*, q^*, z^*)|\theta_1\right](x, q, a)$$

after choosing an action according to

$$\max_a \mathsf{T}^2 \left\{U(x, a) + \mathsf{T}^1 \left[\beta \check{W}^*(y^*, q^*, z^*)|\theta_1\right]|\theta_2\right\}(y, q, a),$$

Recursive Game II

Next consider an approach in which the value function depends only on the reduced information encoded in y, q:

$$W(y,q) = \max_a \mathsf{T}^2 \left(U(x,a) + \mathsf{T}^1 \left[\beta W^*(y^*, q^*) | \theta_1 \right] | \theta_2 \right) (y,q,a)$$

Issues and Extensions

- Incompatible probability assignments over hidden states;
- Construct well defined worst case probabilities over signals - objects that restrict actions, contracts etc.
- Risk-base approach - relax the reduction of compound lotteries as in Segal (1990), Klibanoff, Marinacci, and Mukerji (2003) and Ergin and Gul (2004).
- Constraints instead of penalties - Epstein and Schneider (2003b).
- Penalties that depend on states - Maenhout (2004) and Lin, Pan, and Wang (2004).