Weathering Uncertainty in the Long Run

Lars Peter Hansen
and friends
Thomas J. Sargent, Jose Scheinkman, John Heaton

University of Chicago

Game plan

Connect three topics

1. Recursive and robust decision making with learning - decision theory
2. Long run risk - operator theory
3. Estimation and specification uncertainty in time series - statistical methods

Recursive utility model

CES recursion as in Kreps and Porteus (1978):

$$V_t = \left[(1-\beta)\left(C_t\right)^{1-\rho} + \beta \mathcal{R}_t(V_{t+1})^{1-\rho}\right]^{\frac{1}{1-\rho}}.$$

The recursion incorporates the current period consumption C_t and makes a risk adjustment $\mathcal{R}_t(V_{t+1})$ to the date $t+1$ continuation value V_{t+1} where:

$$\mathcal{R}_t(V_{t+1}) \doteq \left[E\left(V_{t+1}\right)^{1-\alpha} | \mathcal{X}_t\right]^{\frac{1}{1-\alpha}}$$

and $\mathcal{X}_t$ is the current period information set.
Do not reduce compound intertemporal lotteries.
Intertemporal composition of risk matters!

Extraction of macro risk

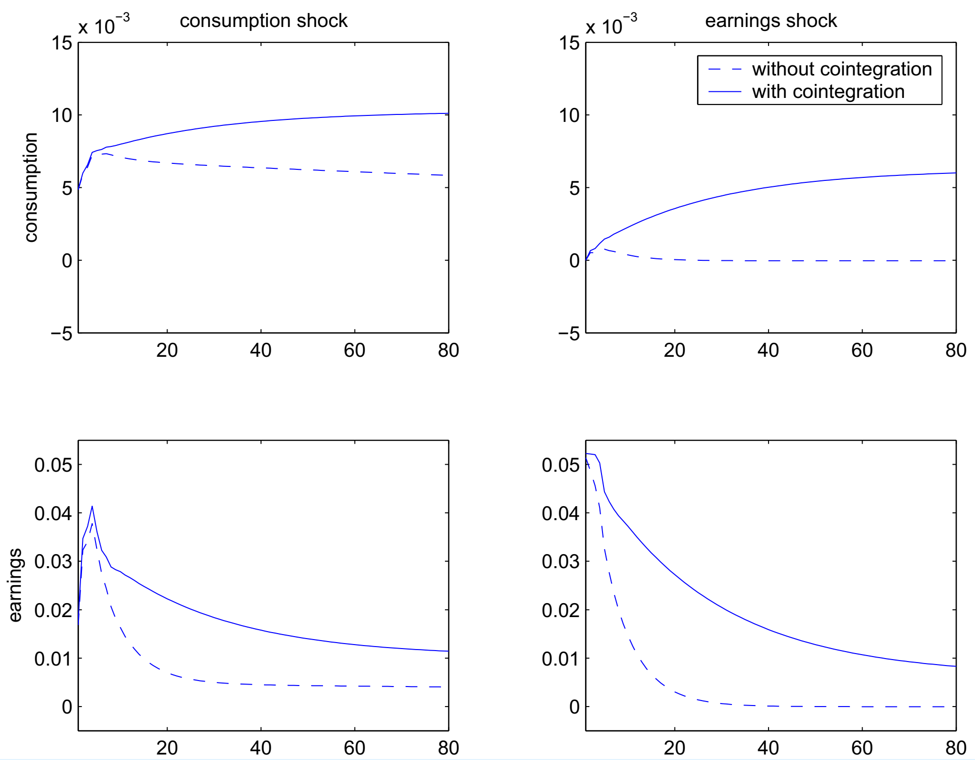

Estimated cash flow responses to macro shocks

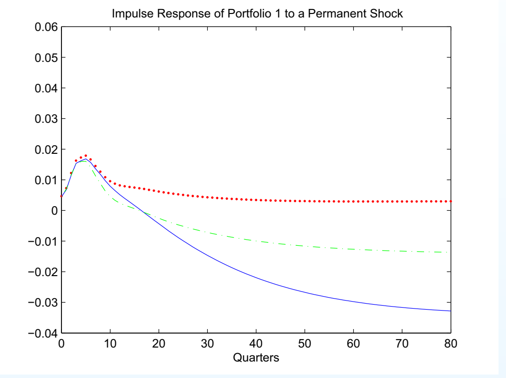

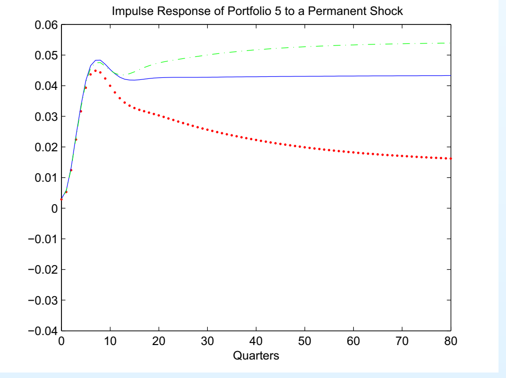

Related modelling challenges

1. Stochastic growth models with hidden growth regimes: productivity slow downs and new economies
2. Permanent income models with uncertain and highly persistent income processes

Issues

- Econometricians face considerable specification uncertainty - especially true in the *long run* - shouldn't economic agents share this burden?
- Probability models are adopted in part for tractability - confront possible mistakes?
- Robust learning and actions should be reflected in forward looking capital values

Game plan for remainder of the talk

1. Review stochastic formulations of robust control problems and explore limitations.
2. Extend the analysis to address some of these issues using hidden state Markov chains. Learning and model averaging.
3. Relate to risk-based counterparts that relax the reduction of compound lotteries.

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, $E(M_t | \mathcal{X}_0) = 1$.
- Distorted probability

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

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

Parameterize alternative probability models via martingales

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

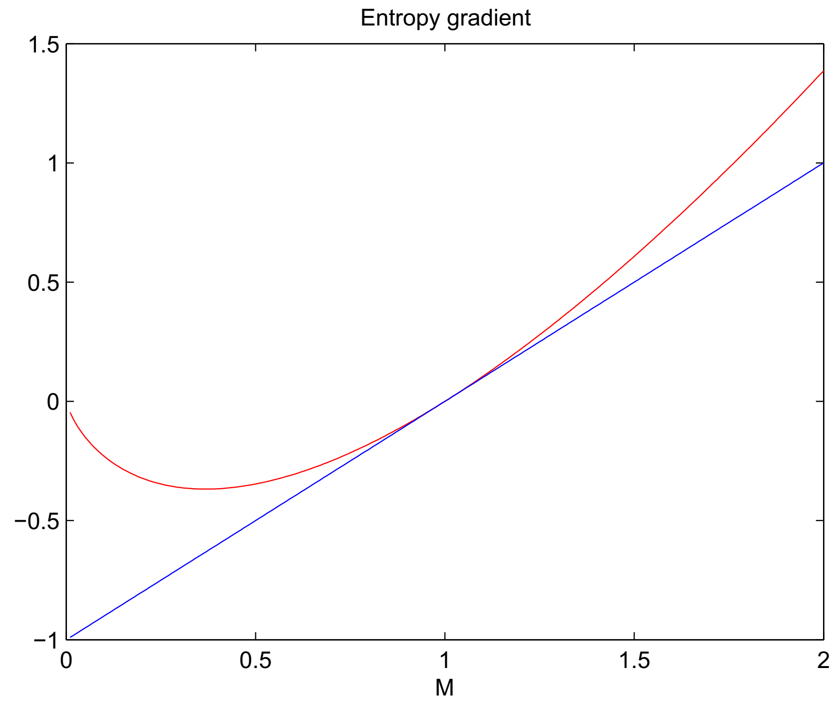

Static entropy penalization problem

- 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

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

- Special case of

$$h^{-1}\left[Eh(V)\right].$$

Static formulation of Robust Control

- Let a be an action in a feasible set $\mathcal{A}$ and x an unknown state. Informational restrictions may be imposed on the action. Actions and states can be processes. Objective can be discounted utility.
- Problem 1

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

  Worst case M^* depends on action. Zero sum game.
- Problem 2 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'.

Dynamic Stochastic Versions of Robust Control

- Petersen, James, and Dupuis (2000) *IEEE Transactions in Automatic Control*
- Anderson, Hansen, and Sargent (2003) European Economic Review
- Hansen, Sargent, Turmuhambetova, and Williams (2004) forthcoming in *JET*

Multiplicative decomposition of a martingale

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 $t+1$.

Factor a joint density as a product of conditionals

Discounted Entropy

- Use a geometric average:

$$(1-\beta)\sum_{j=0}^{\infty}\beta^j E\left[M_{j+1}\log(M_{j+1})|\mathcal{X}_0\right]$$

- Summation by parts:

$$\begin{aligned}(1-\beta) \quad & \textstyle\sum_{j=0}^{\infty} \quad \beta^j E\left[M_{j+1}\log(M_{j+1})|\mathcal{X}_0\right] \\ = \quad & \textstyle\sum_{j=0}^{\infty} \quad \beta^j E\left(M_j E\left[m_{j+1}\log(m_{j+1})|\mathcal{X}_j\right]|\mathcal{X}_0\right)\end{aligned}$$

Robust Control and Discounted Entropy

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

$$V_t = U_t + \beta E(m_{t+1}V_{t+1}|\mathcal{X}_t) + \theta E(m_{t+1}\log m_{t+1}|\mathcal{X}_t)$$

- Link to recursive utility - V_{t+1} continuation value for a future consumption plan - use a risk adjustment:

$$-\theta \log E\left[\exp\left(-\frac{1}{\theta}V_{t+1}\right)|\mathcal{X}_t\right].$$

- Worst case m_{t+1} is: $\frac{\exp\left(-\frac{1}{\theta}V_{t+1}\right)}{E\left[\exp\left(-\frac{1}{\theta}V_{t+1}\right)|\mathcal{X}_t\right]}$ depends on value function.

Questions

1. What about rational expectations? - small likelihood ratios - statistical detection Anderson, Hansen, and Sargent (2003)
2. What about martingale convergence? - Kolomogorov versus Doob Constraining discounted entropy does not require the existence of an absolutely continuous limiting probability. A limiting probability distribution will exist, however.
3. Why not average? - too many models on the table - we will combine with model averaging over smaller dimensions
4. What about learning? - important limitation that I will now address

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. Allow for robustness in the probabilities used for averaging over the hidden states as in prior sensitivity analysis in Bayesian statistics and econometrics

Observations

1. hidden states can be time invariant and hence index alternative models
2. exploit tools for solving hidden state Markov chain models
3. minimizing agent has an informational advantage
4. recursive

An alternative decomposition of M_t

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

- Decompose:

$$M_t = h_t G_t$$

  where $E(h_t|\mathcal{S}_t) = 1$.

- 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:

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

Review: 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

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

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

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

- $\{z_t : t \geq 0\}$ can evolve according to a hidden state Markov chain. Estimate a moving target with known transition probabilities, but unknown transitions. For example, technology has hidden growth states that must be inferred by investors.
- Evolution equations determine a conditional density $\tau(z_{t+1}, s_{t+1} | x_t, a_t)$ relative to the product measure $\lambda \times \eta$.

Information

- Let $\{\mathcal{S}_t : t \geq 0\}$ denote a filtration, where $\mathcal{S}_t$ is generated by y_0, s_1, ..., s_t.
- We can apply Bayes' rule to τ to deduce a density q_t, relative to the measure λ, for z_t conditioned on information $\mathcal{S}_t$.
- Let $\{\mathcal{X}_t : t \geq 0\}$ be a larger filtration where $\mathcal{X}_t$ is generated by x_0, w_1, w_2, ..., w_t.
- Let A denote a feasible set of actions, which we take to be a Borel set of some finite dimensional Euclidean space, and let $\mathcal{A}_t$ be the set of A-valued random vectors that are $\mathcal{S}_t$ measurable.

Benchmark Decision Problem

Decision problem under incomplete information about the state but complete confidence in the model:

$$\max_{a_t \in \mathcal{A}_t : t \geq 0} E \left[\sum_{t=0}^{\infty} \beta^t U(x_t, a_t) | \mathcal{S}_0 \right],$$

and subject to the intertemporal evolution:

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

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,

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

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

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.
- 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 the signal history $\mathcal{S}_t$.
- 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 $V(y_{t+1}, q_{t+1}, z_{t+1})$,

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

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: $\hat{W}_t = \hat{V}(y_t, q_t, z_t, a_t)$ 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

- Construct well defined worst case probabilities over signals - objects that restrict actions, contracts etc; but not over unverifiable or hidden states.
- Related risk-based approach - relax the reduction of compound lotteries as in Segal (1990), Klibanoff, Marinacci, and Mukerji (2003) and Ergin and Gul (2004). Preferences over hidden state risk or subjective state risk are separated from preferences over risk conditioned on the hidden state.
- Constraints instead of penalties - Epstein and Schneider (2003).
- Penalties that depend on states - Maenhout (2004) and Lin, Pan, and Wang (2004).