Constraint and Multiplier Preferences

LPH and TJS

April 2016

Objects

- A baseline probability model
  - controlled Markov process
- Relative entropy
  - quantifies statistical discrepancy between two models
- A set of consumption plans
  - each plan measurable with respect to history of Markov state
- Sets of probability models surrounding a baseline model

Free parameter

EITHER:

- η_0 (for "constraint preferences")
  - Discounted relative entropy pins down a fixed set of models
- θ (for "multiplier preferences")
  - A fixed penalty parameter on contributions to discounted relative entropy

Alternative preference orderings

- Constraint
  - Parameterized by fixed discounted relative entropy η_0
  - *Fixed* set of models that is independent of consumption plan
  - A Lagrange multiplier $\tilde{\theta}$ on an entropy constraint. It depends on
    - η_0
    - the consumption plan
- Multiplier
  - Parameterized by $\theta \in [\underline{\theta}, +\infty]$
  - Discounted relative entropy and associated set of models are functions of consumption plan and θ

Baseline probability model

- $\mathcal{X}_t$ is sigma-algebra generated by x^t.
- $\bar{X}$ is smallest sigma algebra containing x_t for every $t \geq 0$.
- Q is set of all infinite sequences x^∞.
- probability measure π is defined on $(Q, \bar{X})$.
- $q_0(x)$ is a density for x_0.
- $f(x^*, x)$ is a transition density.
- $\pi(x^t)$ is joint density over x^t generated by (f, q_0).

Utility function and plans

- $\{c(x_t)\}_{t=0}^{\infty}$ is a consumption plan
- $u(c(x_t))$ is a utility function
- $\beta \in (0, 1)$ is a discount factor

Likelihood ratios

- $\{M_t\}_{t=0}^{\infty}$ is a nonnegative martingale with $EM_t|\mathcal{X}_0] = 1$
- M_t is a likelihood ratio for each $t \geq 0$
- Distorted density over x^t is

$$\widehat{\pi}(x^t) = M_t \pi(x^t)$$

- $$M_t = E[M^{\infty}|\mathcal{X}_t]$$

  where $M^{\infty}(x^{\infty})$ is a nonnegative random variable with $E_0 M^{\infty}(x^{\infty}) = 1$

Likelihood ratios

Represent $\{M_t\}_{t=0}^{\infty}$ as

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

or

$$M_t = M_0 \prod_{j=1}^{t} m_j,$$

where m_{t+1} is a likelihood ratio that distorts the density f of x_{t+1} conditional on $\mathcal{X}_t$, and where

$$\int m_{t+1}(x_{t+1}|\mathcal{X}_t) f(x_{t+1}|\mathcal{X}_t) dx_{t+1} = 1$$

Decomposition of entropy

$$E_0(M^\infty \log M^\infty) = \sum_{t=0}^{\infty} E_0[M_t E_t(m_{t+1} \log m_{t+1})],$$

where E_t denotes an expectation conditional on $\mathcal{X}_t$.

Discounted entropy

- $$\sum_{t=0}^{\infty} \rho^{t+1} E_0[M_t E_t(m_{t+1} \log m_{t+1})], \quad \rho \in (0,1)$$
- We'll set $\rho = \beta$

Discounting?

Choice: what to discount and at what rates

- Period utilities?
- Multiplicative entropy contributions m_{t+1}?

We discount both with common fixed discount factor β

Discounting?

When m_{t+1} is not discounted, the minimizing agent exhausts his entropy budget early

- that makes
$$\lim_{t \to +\infty} E_t(m_{t+1} \log m_{t+1}) = 0,$$
so that concerns about robustness die at a rate that depends on β.

Constraint preferences

$$v(x_0, \eta_0) = \min_{\{m_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \int u(c(x_t)) M_t(x^t) \pi(x^t) dx^t$$

subject to

$$\begin{aligned}
M_{t+1}(x^{t+1}) &= m(x_{t+1}|x^t) M(x^t) \\
\int m(x_{t+1}|x^t) f(x_{t+1}|x_t) dx_{t+1} &= 1, \quad t \geq 0 \\
\sum_{t=0}^{\infty} \beta^{t+1} E_0[M_t E_t(m_{t+1} \log m_{t+1})] &\leq \eta_0 \\
M_0 &= 1
\end{aligned}$$

Constraint preferences

$$v(x_0, \eta_0) = \min_{\{m_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \int u(c(x_t)) M_t(x^t) \pi(x^t) dx^t$$

subject to

$$\begin{aligned}
M_{t+1}(x^{t+1}) &= m(x_{t+1}|x^t) M(x^t) \\
\int m(x_{t+1}|x^t) f(x_{t+1}|x_t) dx_{t+1} &= 1, \quad t \geq 0 \\
\sum_{t=0}^{\infty} \beta^{t+1} E_0[M_t E_t(m_{t+1} \log m_{t+1})] &\leq \eta_0 \quad : \tilde{\theta} \\
M_0 &= 1
\end{aligned}$$

Lagrangian for constraint preferences

$$\min_{\{m_{t+1}\}_{t=0}^{\infty}} \max_{\tilde{\theta}} E_0 \sum_{t=0}^{\infty} M_t \beta^t \left[u(c(x_t)) + \beta \tilde{\theta} m_{t+1} \log m_{t+1} \right] - \tilde{\theta} \eta_0$$

subject to

$$\begin{aligned} M_{t+1}(x^{t+1}) &= m(x_{t+1}|x^t) M(x^t) \\ M_0 &= 1 \\ E_t m_{t+1} &= 1 \forall t \geq 0 \end{aligned}$$

Multiplier Preferences

$$w(x_0) = \min_{\{m_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} M_t \beta^t \left[u(c(x_t)) + \beta\theta m_{t+1} \log m_{t+1} \right]$$

subject to

$$\begin{aligned} M_{t+1}(x^{t+1}) &= m(x_{t+1}|x^t) M(x^t) \\ M_0 &= 1 \\ E_t m_{t+1} &= 1 \forall t \geq 0 \end{aligned}$$

Indirect utility function

$$
\begin{aligned}
\mathsf{T}(w(x^*))(x) &= -\theta \log \int \exp\left(\frac{-w(x^*)}{\theta}\right) f(x^*|x)dx^* \\
&= \min_{m(x^*,x)} \int m(x^*,x)[w(x^*) + \theta \log m(x^*,x)] f(x^*|x)dx^*
\end{aligned}
$$

subject to $\int m(x^*,x) = \int m(x^*,x) f(x^*|x)dx^* = 1$

The minimizing m is

$$
\check{m}(x^*,x) \propto \exp\left(\frac{-w^*(x^*)}{\theta}\right)
$$

Indirect utility function

$$
\begin{aligned}
\mathsf{T}(w(x^*))(x) &= -\theta \log \int \exp\left(\frac{-w(x^*)}{\theta}\right) f(x^*|x)dx^* \\
&= \min_{m(x^*,x)} \int m(x^*,x)[w(x^*) + \theta \log m(x^*,x)] f(x^*|x)dx^*
\end{aligned}
$$

subject to $\int m(x^*,x) = \int m(x^*,x) f(x^*|x)dx^* = 1$

The minimizing m is

$$
\check{m}(x^*,x) \propto \exp\left(\frac{-w^*(x^*)}{\theta}\right)
$$

- Statistical version of Murphy's law
- Easy to compute

Recursive version of multplier problem

Value function

$$w(x) = u(c(x)) + \beta \mathsf{T}(w(x^*))(x)$$

Associated minimizing likelihood ratio

$$\check{m}(x^*, x) \propto \exp\left(\frac{-w^*(x^*)}{\theta}\right)$$

Recursive representation of discounted entropy

Recursive representation:

$$R_0 = \sum_{t=0}^{\infty} \beta^{t+1} E_0[M_t E_t(m_{t+1} \log m_{t+1})]$$

$$R_0 = \sum_{\tau=1}^{\infty} \beta^{\tau} E_0\left[\frac{M_{\tau-1}}{M_0} E_{\tau-1}(m_{t+\tau} \log m_{t+\tau})\right]$$

$$R_t = \sum_{\tau=1}^{\infty} \beta^{\tau} E_t\left[\frac{M_{t+\tau-1}}{M_t}(m_{t+\tau} \log m_{t+\tau})\right]$$

$$M_t R_t = E_t \sum_{\tau=1}^{\infty} \beta^{\tau} M_{t+\tau-1} E_{t+\tau-1}(m_{t+\tau} \log m_{t+\tau}).$$

Recursive representation of discounted entropy

$$R_t = \beta E_t \big(m_{t+1} [\log m_{t+1} + R_{t+1}] \big)$$

Quiz

Does

$$R_t = \beta E_t \big(m_{t+1}[\log m_{t+1} + R_{t+1}] \big)$$

remind you of a "promise keeping constraint"?

Recursive Represenation of Constraint Preferences

$$v(x, R) = u(c(x)) + \min_{m, R^*} \beta \int m(x^*, x) v(x^*, R^*) f(x^*|x) dx^*$$

subject to

$$\begin{aligned} \beta \int m(x^*, x) \left[\log m(x^*, x) + R^*(x^*) \right] f(x^*|x) dx^* &\leq R \\ \int m(x^*, x) f(x^*|x) dx^* &= 1 \end{aligned}$$

State variable degeneracy

$$v_R(x, R) = -\tilde{\theta}$$

Staring implies that R is a function of x

Exploiting connections

- Constraint problem
  - The one we are really interested in, but ...
  - It is difficult to solve directly
- Multiplier problem
  - Less interesting than constraint problem, but ...
  - Solving it is a useful intermediate step in solving the constraint problem

Magic three step process

- Fix θ and solve Bellman equation for multiplier value function $w(x)$ and associated likelihood ratio $m(x^*, x)$.
- Compute implied relative entropy by solving this functional equation for $R(x)$:

$$R(x) = \beta \int (m(x^*, x)[\log m(x^*, x) + R(x^*)]f(x^*, x)dx^*$$

- Check whether $R(x_0) = \eta$
  - if yes, then

    $$v(x_0, \eta) = w(x) + \beta R(x_0)$$

    and

    $$\tilde{\theta} = \theta$$

  - If not, try another θ.

Why constraint preferences interest us more

- Key parameter of constraint preferences is R_0
  - R_0 is closely linked to statistical measures for discriminating between statistical models
  - This gives a way to pick plausible values of R_0
  - If we compare alternative consumption plans $c(x_t)$, it is natural to hold R_0 fixed.

Why constraint preferences interest us more

- Key parameter of multiplier preferences is θ
  - There is no natural way to quantify θ without consulting the *context*, meaning the utility function, the set of consumption plans under consideration, and the baseline transition density $f(x^*, x)$.
  - A fixed θ can imply very different values of R_0 different contexts.

Two Preference Orderings

- $v(x, \eta_0)$ and $w(x; \theta)$ induce *different* preferences over alternative plans $c(x)$.
- It is easier to compute $w(x)$ than $v(x, R)$.
- Fixed entropy ball R versus fixed "penalty parameter" θ