Model Uncertainty and Robust Control in Monetary Policy

Noah Williams

University of Wisconsin - Madison

November 23, 2015

Introduction

- "Uncertainty is not just an important feature of the monetary policy landscape; it is the defining characteristic of that landscape." Greenspan (2003).
- Uncertainty is pervasive. Some sources: models, parameters, shocks, data quality.
- Policymakers must continually face such uncertainty. Structural models of dynamic economies require specifications of how people respond to risk and uncertainty.
- Typical approach assumes people confront risk using a model of rational expectations. Decision-makers are endowed with the correct model to use in assessing risk. No scope for modelling error or approximation.

Model Uncertainty

- Consider methods which account for model uncertainty.
- Specify sets of possible models, how to incorporate them in decisions and policy choice.
- Develop policies which are robust to model uncertainty, in that they perform well across a range of alternative models.
- Policy choices may be sensitive to assumptions about uncertainty. Rules "robust" to one type of uncertainty may perform poorly when faced with uncertainty of a different type.

Model Uncertainty: Definitions

- Consider the design of decision rules that perform well across a range of alternative models.
- A **model** is a specification of a probability distribution over outcomes of interest to the decision maker, which is influenced by a control variable.
- **Model uncertainty**: decision maker faces subjective uncertainty about the specification of this probability distribution.

Model Uncertainty: Issues

- **First issue:** specify the class of alternative models which the decision maker entertains.
- Often take a benchmark nominal model and consider perturbations of this model. How to specify and measure the magnitude of the perturbations.
- **Second issue:** choose a decision rule and thus what it means for a rule to “perform well” across models.
- Main approaches:
  - Bayesian: decision maker forms a prior over models and maximizes expected utility. No real difference between model uncertainty or uncertainty over shocks.
  - Robust: minimize the worst case loss over the set of possible models. Distinguish between shocks which are averaged over, and models which are not. Decision makers are either unable or unwilling to form a prior over the forms of model misspecification, but can bound the class of alternative models.

Example: Parameter Uncertainty like Brainard (1967)

- Consider model along the lines of Kydland-Prescott/Barro-Gordon:

$$U = U^* - \gamma(\pi - \pi^e) + \varepsilon \tag{1}$$

- Policymakers choose π, π^e given, $\varepsilon \sim N(0,1)$. For simplicity assume $\pi^e = 0$.
- Policy problem, no model uncertainty:

$$\begin{aligned} & \min_{\pi} E[U^2 + \pi^2] \text{ s.t. } (1) \\ = \quad & \min_{\pi} E[(U^*)^2 - 2\gamma U^* \pi + \gamma^2 \pi^2 - \gamma \pi \varepsilon + \varepsilon^2 + \pi^2] \end{aligned}$$

- Optimal policy:

$$\pi = \frac{\gamma}{1+\gamma^2} U^*$$

- Uncertainty through ε has no effect on policy.

Parameter Uncertainty

- Now suppose that slope of Phillips curve γ is unknown.
- Uncertainty about it is captured by a prior distribution: $\gamma \sim N(\bar{\gamma}, \sigma^2)$ independent of ε.
- Now expectation is over ε and γ:

$$\min_{\pi} E[(U^*)^2 - 2\gamma U^* \pi + \gamma^2 \pi^2 - \gamma \pi \varepsilon + \varepsilon^2 + \pi^2]$$

- Optimal policy:

$$\pi = \frac{E(\gamma)}{E(1+\gamma^2)} U^* = \frac{\bar{\gamma}}{1+\bar{\gamma}^2+\sigma^2} U^*$$

- Uncertainty about γ leads to **caution**, response coefficient is smaller.

Parameter Uncertainty: Implications

- Many policymakers seem to have internalized the lesson of the example, intuitively argue uncertainty leads to more cautious policy. Ex.: Blinder (1998)
- Note however that the result is special. If instead of being independent of ε, the slope γ is correlated with it:

$$E[\gamma\varepsilon] = \rho \neq 0$$

- Now term in $\gamma\varepsilon$ does not vanish:

$$\min_{\pi} E[(U^*)^2 - 2\gamma U^*\pi + \gamma^2\pi^2 - \gamma\pi\varepsilon + \varepsilon^2 + \pi^2]$$

- Optimal policy:

$$\pi = \frac{E(\gamma U^* - \gamma\varepsilon)}{E(1+\gamma^2)} = \frac{\bar{\gamma} U^* - \rho}{1 + \bar{\gamma}^2 + \sigma^2}$$

- If $\rho < 0$ then the response coefficient could be larger. Bigger shocks ε imply smaller response of U to π, so π must respond more aggressively.

Ambiguity or Uncertainty Aversion

- Ellsberg (1961) paradox:
  2 urns, 100 balls, some red some black
  Urn A: 50 red, 50 black, Urn B: unknown
  Bet 1: ball from urn A is black (AB)
  Bet 2: ball from urn A is red (AR)
  Similarly for BB, BR.
  Observed empirically: AB $\sim$ AR $>$ BB $\sim$ BR
- Gilboa and Schmeidler (1989):
  For urn B, subject has too little information to form a prior. Considers a **set** of priors as possible. Being uncertainty averse, evaluates by **minimal** expected utility over priors in the set.

Robustness

- Burl, *Linear Optimal Control* (1999): The control system engineer should be assured that a design will function acceptably before committing to implementation. Such assurance can be obtained by analyzing control system stability and performance with respect to a **range of plant models** that is expected to encompass the actual plant. This type of analysis is termed robustness analysis.
- Huber (1981): ... as we defined robustness to mean insensitivity with regard to small deviations from assumptions, any quantitative measure of robustness must somehow be concerned with the maximum degradation of performance possible for an ϵ-deviation from the assumptions. The optimally robust procedure minimizes this degradation and hence will be a **minimax** procedure of some kind.

Robustness and Ambiguity Aversion

- Robust control started in control theory literature in 1980s, widely used now.
- Linear quadratic Gaussian control widespread in 1970s. For models with unobserved states, performance of rules not robust to small perturbations.
- Influential example: Doyle (1978)
  Title: "Guaranteed Margins for LQG Regulators"
  Abstract: "There are none."
- Hansen and Sargent (and co-authors) have shown that robust control can be viewed as an implementation of ambiguity aversion.
- **Key:** "Model" = probability measure
  set of priors = class of models/measures
  Lagrange multiplier theorem links formulations

Brainard under Robustness

- Now suppose that there may be more general uncertainty about the model, which we incorporate as an additional shock.
$$U = U^* - \gamma(\pi - \pi^e) + \varepsilon + w \tag{2}$$
- Alternative models are indexed by w. Approximating or nominal model sets $w = 0$.
- Consider models close to the approximating model, with log-likelihood ratios or relative entropy bounded by η.
- Since $\varepsilon \sim N(0, 1)$ the bound is $w^2 \leq \eta$.
- Robust control problem, constraint version:
$$\min_{\pi} \max_{w} E[U^2 + \pi^2] \text{ s.t. } (2),\ \ w^2 \leq \eta$$
- A particular implementation of Gilboa-Schmeidler, puts structure on set of priors

Robust Multiplier Problem

- Robust control problem, multiplier version:

$$\min_{\pi} \max_{w} E[U^2 + \pi^2] + \theta w^2 \text{ s.t. } (2)$$

- Via Lagrange multiplier theorem, can link constraint and multiplier problems. θ is multiplier on the constraint.
- Multiplier model easier to work with in practice, of interest in its own right Strzalecki (2011).
- Solutions:

$$\begin{aligned}
\pi(\theta) &= \frac{\gamma\theta}{\theta - 1 + \gamma^2\theta}(U^* + \gamma\pi^e) \\
w(\theta) &= \frac{1}{\theta - 1 + \gamma^2\theta}(U^* + \gamma\pi^e) \\
\eta &= w^2 = \left(\frac{1}{\theta - 1 + \gamma^2\theta}\right)^2 (U^* + \gamma\pi^e)^2
\end{aligned}$$

- Inverse relationship between θ and η. As $\theta \to \infty$ then $\eta \to 0$ and we are back at problem w/o robustness.

Performance of Robust Rules

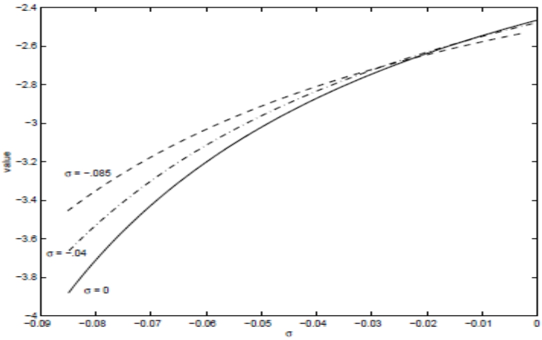

- Decision rules as $\sigma = -1/\theta$, x-axis assumes $w(\sigma_2)$ while policy is $\pi(\sigma_1)$.
- Non-robust rule ($\sigma_1 = 0$) best when no misspecification ($\sigma_2 = 0$) but performance deteriorates when misspecification increases ($\sigma_2 < 0$).
- Robust rules sacrifice performance under nominal model for better performance under misspecification.

Robustness: Dynamic Version

- Now extend these static considerations to a dynamic environment, as in Hansen and Sargent.
- Ideas extend more generally, but focus on LQ model.
- Nominal model:

$$x_{t+1} = Ax_t + Bi_t + C\varepsilon_{t+1}, \tag{3}$$

  x_t state, i_t agent's control, ε_t is an i.i.d. $N(0, I)$
- Agent's intertemporal preferences:

$$E_0 \sum_{t=0}^{\infty} \beta^t \left(x_t' Q x_t + i_t' R i_t \right) \tag{4}$$

  $0 < \beta < 1$, Q and R are negative definite.
- Perturb nominal model with an "misspecfication shock" w_{t+1}, allowed to be correlated with state x_t:

$$x_{t+1} = Ax_t + Bi_t + C(\varepsilon_{t+1} + w_{t+1}). \tag{5}$$

Set of Alternative Models

- As before w_{t+1} used to represent alternative models which are made to close to the nominal model in an entropy or log-likelihood sense.
- Nominal and distorted model distributions:

$$\begin{aligned} f_0(x'|x) &\sim N(Ax + Bi, CC') \\ f(x'|x) &\sim N(Ax + Bi + Cw, CC') \end{aligned}$$

- Relative entropy is defined as expected log-likelihood ratio:

$$\begin{aligned} I(f_0, f) &= \int \log\left(\frac{f(y'|y)}{f_0(y'|y)}\right) f(y'|y)\, dy' \\ &= \frac{1}{2} w' w \end{aligned}$$

- For dynamic model, consider discounted sum of one-step entropies:

$$E_0 \sum_{t=0}^{\infty} \beta^t w'_{t+1} w_{t+1} \leq \eta \tag{6}$$

Dynamic Robust Multiplier Problem

- Agent maximizes (4) with respect to the worst case perturbed model (5) from the set (6).
- Using a Lagrange multiplier theorem, the constraint set can be converted to a penalty and the decision problem can by solved recursively by solving the Bellman equation for a two player zero sum game:

$$V(x) = \max_i \min_w \left\{ x'Qx + i'Ri + \beta\theta w'w + \beta E\left[V(x')|x \right] \right\}$$

  s.t. (5), $\theta > 0$ Lagrange multiplier on (6) and the expectation over ε.
- As $\theta \to \infty$ only the nominal model remains (thus $\eta \to 0$), and the decision rule is less robust.
- Conversely, there is typically a minimal value of θ beyond which the value is $V(x) = -\infty$. This gives the most robust decision rules, allowing for the largest uncertainty set.

Dynamic Robust Multiplier Problem: Solution

- Given the LQ nature of the problem, the solution of the robust multiplier problem is a pair of linear decision rules for i and w, and quadratic value function:

$$i_t = -Fx_t, \quad w_{t+1} = Kx_t, \quad V(x_t) = -x_t' P x_t$$

- Solution can be found by considering the minimizing problem taking as given F:

$$\min_w \theta w' w - [(A - BF)x + Cw]' P [(A - BF)x + Cw]$$

- Solution is:

$$w = \theta^{-1}(I - \theta^{-1} C' P C)^{-1} C' P (A - BF)x$$

with value:

$$-x'(A - BF)' \mathcal{D}(P)(A - BF)x$$

$$\text{where:} \qquad \mathcal{D}(P) = P + PC(\theta I - C' P C)^{-1} C' P$$

Dynamic Robust Multiplier Problem: Solution

- Then given this solution, consider maximizing problem:

$$\max_i x'Qx + i'Ri - \beta(Ax+Bi)'\mathcal{D}(P)(Ax+Bi)$$

- Solution of this is like a usual Riccati equation, but with $\mathcal{D}(P)$ as continuation.

$$\begin{aligned} i &= -\mathcal{F} \circ \mathcal{D}(P)x \\ \mathcal{F}(\Omega) &= \beta[R + \beta B'\Omega B]^{-1} B'\Omega A \end{aligned}$$

with value:

$$-x' T \circ \mathcal{D}(P)x$$

where: $T(P) \;=\; Q + \beta A'(P - \beta PB(R + \beta B'PB)^{-1}B'P)A$

- Value matrix P solves the modified Riccati equation:

$$P = T \circ \mathcal{D}(P)$$

Application: Robust Monetary Policy

- Onatski and Stock (2002) consider robust Taylor rules in a backward-looking monetary model of Rudebusch and Svensson (1999):

$$\begin{aligned}
\pi_{t+1} &= a_0\pi_t + a_1\pi_{t-1} + a_2\pi_{t-2} + (1 - a_0 - a_1 - a_2)\pi_{t-3} \\
&\quad + a_y y_t + \varepsilon_{\pi,t+1} \\
y_{t+1} &= b_0 y_t + b_1 y_{t-1} - b_r(\bar{i}_t - \bar{\pi}_t) + \varepsilon_{y,t+1}
\end{aligned}$$

  where π_t=inflation, y_t=output gap, and $\bar{\cdot}$=annual average
- Consider policy rules of the class: $i_t = g_\pi \bar{\pi}_t + g_y y_t$
- Found optimal rule w/o robustness more aggressive than Taylor rule, robust policy rules even more aggressive

Robust Optimal Taylor Rules

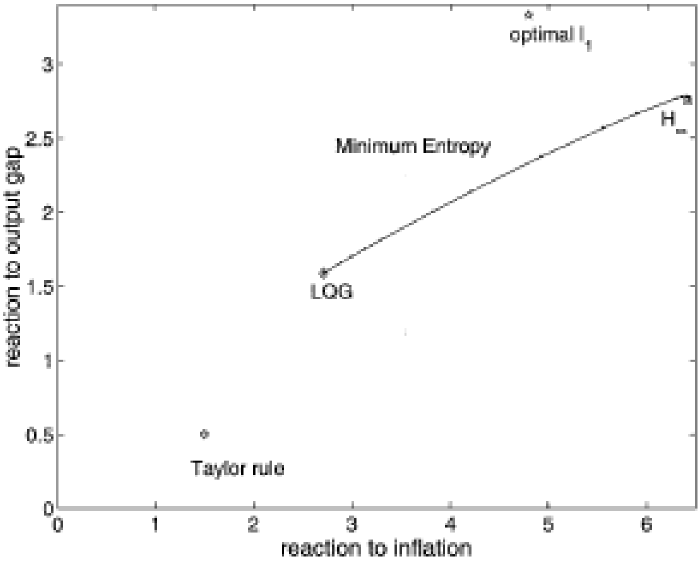

Loss Function in the Most Robust Case (minimal θ)

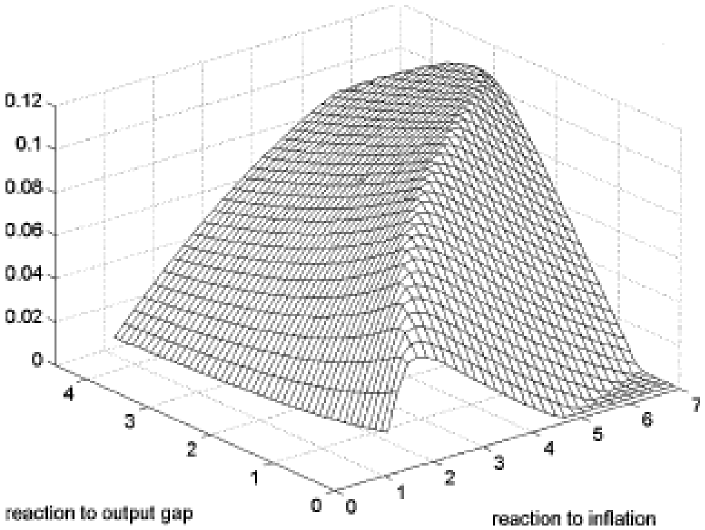

Application: Robust Monetary Policy under Discretion

- Leitemo and Soderstrom (2008) consider robust optimal monetary policy in the New Keynesian model without commitment.

$$\begin{aligned}\pi_t &= \beta E_t \pi_{t+1} + \kappa x_t + C\pi[w_t^\pi + \varepsilon_t^\pi] \\ x_t &= E_t x_{t+1} - \sigma^{-1}(i_t - E_t \pi_{t+1}) + C_x[w_t^x + \varepsilon_t^x]\end{aligned}$$

- Period loss function $\pi_t^2 + \lambda x_t^2$. Find optimality conditions:

$$\begin{aligned}x_t &= -\frac{\kappa}{\lambda}\pi_t \\ w_t^\pi &= \frac{C_\pi}{\theta}\pi_t \\ w_t^x &= 0\end{aligned}$$

Implications

- Solving for equilibrium dynamics, obtain:

$$i_t = \frac{\sigma\kappa}{\lambda(1 - C_\pi^2/\theta) + \kappa^2} C_\pi \varepsilon_t^\pi + \sigma C_x \varepsilon_t^x$$

- The central bank responds to a positive cost shock by tightening policy (as under the nonrobust policy).
- But coefficient decreasing in θ, so more robustness (smaller θ) leads to more aggressive response to cost shocks.
- Central bank fears cost shocks have larger impact on inflation, so output and interest rates more volatile, but inflation less volatile.

Application: Robust Monetary Policy with Near-Rational Expectations

- Wooford (2010) considers model where central bank trusts its model of the economy, but is concerned about private sector misspecifications.
- Private agents behavior consistent with aggregate supply:

$$\pi_t = \kappa x_t + \beta E_t[m_{t+1}\pi_{t+1}] + u_t$$

  where m_{t+1} is the belief distortion. Policymakers believe $u_t = \sigma_u \varepsilon_t$, but not private agents.
- Policymakers solve the problem:

$$\min_{\{\pi_t\}} \max_{\{m_{t+1}\}} E_{-1} \sum_{t=0}^{\infty} \beta^t [\pi_t^2 + \lambda x_t^2] - \theta E_{-1} \sum_{t=0}^{\infty} \beta^t m_{t+1} \log m_{t+1}$$

- Note policymaker expectations not distorted, but consider any private sector beliefs satisfying the bound.

Results

Woodford shows optimal policy can be written:

$$\pi_t = \mu\pi_{t-1} + \bar{p}^1\varepsilon_t - \mu\sigma_u\varepsilon_{t-1}$$

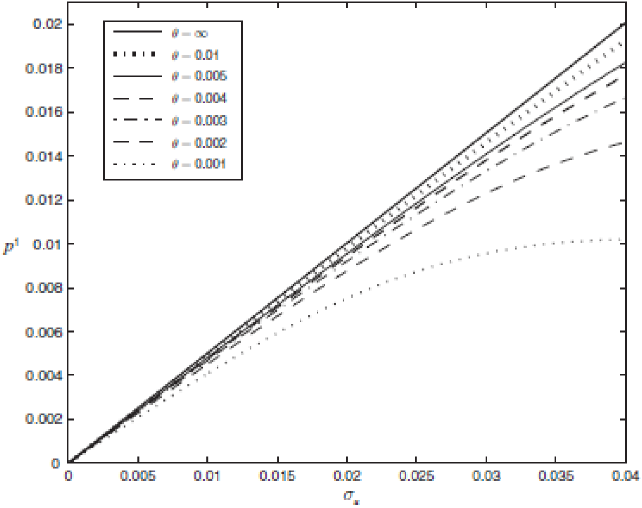

Impulse Responses to Cost-Push Shock

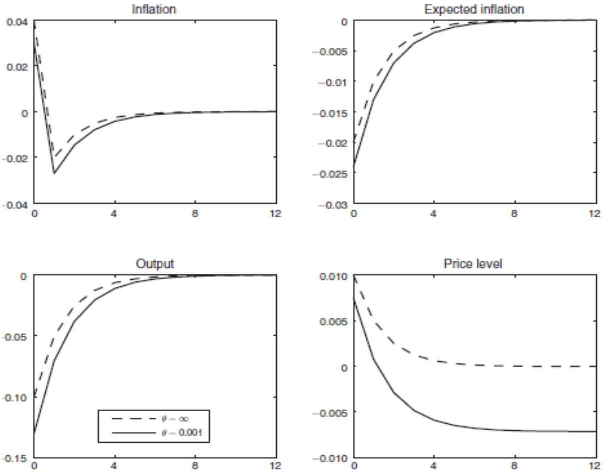

Robust optimal policy more history dependent, reduces price below level would have prevailed in absence of shock

Application: Robust Asset Pricing

- Although we focus on monetary policy in this lecture, much of the literature on robustness has focused on implications for asset prices.
- The equity premium puzzle suggests that agents act very risk averse in asset markets. This literature studies whether a small degree of robustness or ambiguity aversion can substitute for a large amount of risk aversion.
- Hansen, Sargent, and Tallarini (1999), Cagetti, Hansen, Sargent, and Williams (2002) show that robustness adds an additional motive for precautionary saving. This leads to more capital accumulation.
- However making agents more impatient at same time they are more robust can preserve quantity implications.
- Risk prices however are increased under robustness. "Reasonable" amounts of robustness help to solve equity premium puzzle but cannot account for all of it.

Robust Asset Pricing with Unobserved States

- Cagetti, Hansen, Sargent, and Williams (2002): environment with partial information, jumps
- Analyze stochastic growth model with unobserved growth rate. Filter data to predict state.
- Robustness adds additional motive for precautionary saving, can be offset with impatience.
- Robustness adds to the market price of risk: small model uncertainty helps solve equity premium puzzle.
- Time variation in MPR: highest going in and out of recessions. Robustness lowers P/E ratios, but can't explain volatility.

Technology Process

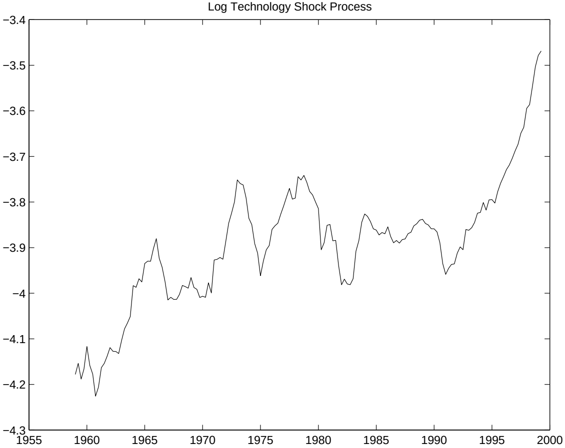

Probability of Being in Low Growth State

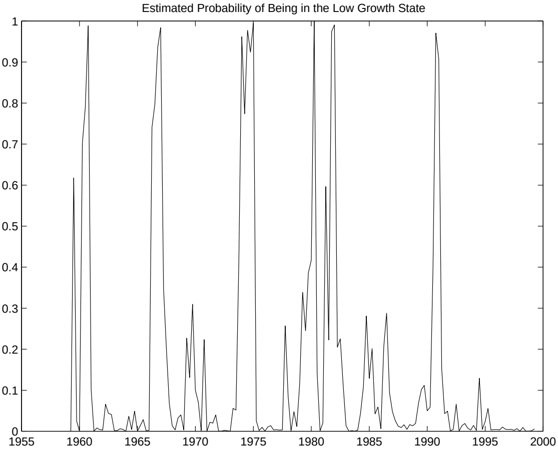

Capital Accumulation Under Robustness

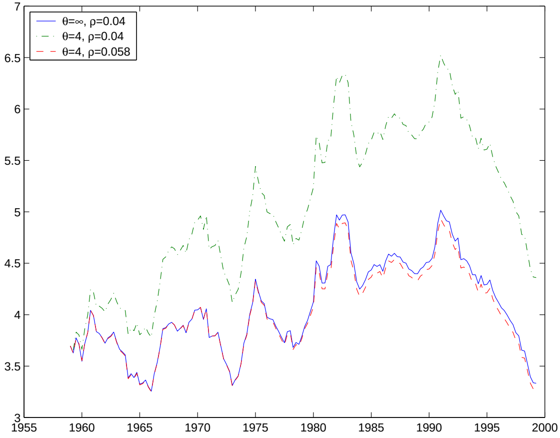

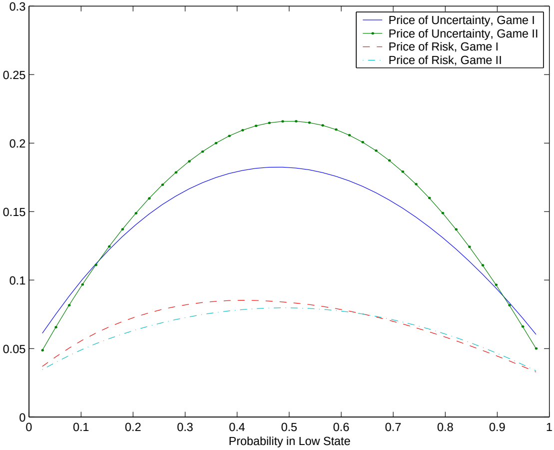
0.3
0.25
0.2
0.15
0.1
0.05
0
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
Probability in Low State
Price of Uncertainty, Game I
Price of Uncertainty, Game II
Price of Risk, Game I
Price of Risk, Game II

Adding Structure to the Uncertainty

- So far the model uncertainty has been unstructured: all models that satisfy entropy bound.
- However some of the models close to the nominal model in a statistical sense may not be plausible economically.
- Decision makers may have a discrete set of models in mind, and bounding them all in one uncertainty set may include extraneous implausible models.
- Decision makers may be more confident some aspects of the model relative to others.

Example: Adding Structure

- Consider the same nominal model (3) as above, but suppose uncertainty is instead solely in parameters A and B:
$$x_{t+1} = (A + \hat{A})x_t + (B + \hat{B})i_t + C\varepsilon_{t+1} \tag{7}$$
for some matrices $\hat{A}$ and $\hat{B}$.
- Can write (7) as a version of (5) with:
$$w_{t+1} = \hat{A}x_t + \hat{B}i_t, \tag{8}$$
- So parametric perturbations a special case of the unstructured uncertainty.
- But how uncertainty is measured matters: whether bound w_{t+1} in entropy sense or restrict $(\hat{A}, \hat{B})$ to confidence intervals. Measurements also will depend on actual control rule for i_t in place.

Uncertainty Models: An Example

- Onatski and Williams (2003) show structure of uncertainty matters, as does measure of "size" of model set.
- Again Rudebusch-Svensson (1999) nominal model:

$$\begin{aligned}
\pi_{t+1} &= \underset{(.08)}{.70}\,\pi_t - \underset{(.10)}{.10}\,\pi_{t-1} + \underset{(.10)}{.28}\,\pi_{t-2} + \underset{(.08)}{.12}\,\pi_{t-3} \\
&\quad + \underset{(.03)}{.14}\,y_t + \varepsilon_{\pi,t+1} \\
y_{t+1} &= \underset{(.08)}{1.16} y_t - \underset{(.08)}{.25}\,y_{t-1} - \underset{(.03)}{.10}\,(\bar{\imath}_t - \bar{\pi}_t) + \varepsilon_{y,t+1}
\end{aligned}$$

- Consider rules and loss of the form:

$$\begin{aligned}
i_t &= g_\pi \bar{\pi}_{t-1} + g_y y_{t-2} \\
L_t &= \bar{\pi}_t^2 + y_t^2 + \frac{1}{2}(i_t - i_{t-1})^2.
\end{aligned}$$

- How to model & measure uncertainty: structured vs. unstructured

Example Continued

- (Simple) Structured approach: consider one extra lag of y in Phillips curve, one extra lag of real rate in IS equation. Re-estimate model and consider different parametric models in confidence ellipsoid around point estimates.
- (Simple) Unstructured approach: Consider all model errors w_{1t}, w_{2t} of the Phillips curve, IS that satisfy an entropy bound:

$$E \sum_{t=0}^{\infty} \beta^t \left(\frac{w_{1t}^2}{\text{Var}(\varepsilon_{\pi t})} + \frac{w_{2t}^2}{\text{Var}(\varepsilon_{yt})} \right) < \eta.$$

- As $\beta \to 1$, with η at its maximum value, get H_∞ optimal rule. Turns out to be very aggressive.
- But the "extremely robust" rule destabilizes economy for parametric deviation inside a 20% confidence ellipsoid.

Lesson of the Example

- "Extremely robust" rule performs poorly because not designed for parametric uncertainty. For example, uncertainty about slope of IS curve very influential for aggressive policy rules. Unstructured uncertainty doesn't consider these effects.
- When calibrating uncertainty with Hansen-Sargent method, measured uncertainty is conditional upon a policy rule. Doesn't capture structural uncertainty.
- Robust rules may be fragile for small changes in uncertainty models. When designing policy rules, must **model** model uncertainty.