ECON 2021 - FINANCIAL ECONOMICS I

Lecture 10 – Model Uncertainty & Robust Control

November 19, 2018

MOTIVATION

- Last week we saw that recursive preferences can mitigate the Risk-Free Rate Puzzle by disentangling intertemporal substitution from risk aversion.
- However, a high degree of risk aversion is still needed to explain returns on risky assets.
- The Long-Run Risk model can reduce the required degree of risk aversion when $\gamma > \rho$.
- However, the evidence for predictable consumption growth and $EIS > 1$ is weak.
- Today we study ambiguity aversion and robust control. Two key results emerge:
  1. An apparently high price of risk might partly reflect a price of model uncertainty. This allows us to substitute ambiguity aversion in place of risk aversion, and thus explain asset returns with less risk aversion.
  2. The results of the LRR model can be obtained without the actual presence of long-run risk. It is sufficient that agents suspect that such risk is present.

History

- Knight (1921): There is a difference between risk and uncertainty.
- Keynes (1936): Financial markets are volatile because they price uncertainty, not risk.
- Savage (1954): Under certain axioms, the distinction between risk and uncertainty is behaviorally irrelevant.
- Ellsberg (1961): Experiments cast doubt on the Savage axioms.
- "It takes a model to beat a model". SEU dominates in economics until the 1990s (Gilboa & Schmeidler (1989)). Coincidentally, engineers develop Robust Control methods at around the same time.
- Hansen & Sargent brought these two literatures together.

The Ellsberg Paradox

- Consider 2 urns, each known to contain 100 Red & Black balls.
  1. Urn 1: Known to have 50 Red/50 Black balls.
  2. Urn 2: No info given about Red vs. Black balls.

- Now suppose you can bet on the color of the ball drawn from one of the urns. You get $100 if you guess correctly, $0 otherwise.

- The SEU axioms predict that people should (weakly) prefer Urn 2. According to SEU, people have unique subjective prior probabilities about the proportion of R vs. B in Urn 2. This prior must assign at least a probability of 0.5 to either R or B, so Urn 2 should be (weakly) preferred.

- In practice, the vast majority of subjects prefer Urn 1. The rationale? It is better to wager with ‘known probabilities’.

- Gilboa & Schmeidler (1989) relax the SEU axioms to incorporate a preference for known probabilities. According to SEU, if an individual is indifferent between two lotteries, he should be indifferent to a mixture between them (with known probabilities). The axioms of GS allow agents to (weakly) prefer the mixture.

Model Uncertainty vs. Parameter Uncertainty

- Bayesian Decision Theory does not admit a distinction between Model Uncertainty and Parameter Uncertainty.
- BDT reduces everything to parameter uncertainty. If an agent is unsure about a collection of candidate models, he assigns prior probabilities to them, and then formulates a single 'hypermodel' by averaging over them.
- For example, let $f(y|x,d)$ represent a model for the state, y, tomorrow, given the state, x, and decision, d, today. If an agent entertains multiple candidate models, he can formulate a *family* of models, $g(y|x,d,m)$, indexed by the parameter m.
- According to BDT, he must have a unique prior, $\pi(m)$, over m. Hence, we can reduce the problem to one with a single (subjectively known) model by a process of Bayesian Model Averaging

$$f(y|x,d) = \int g(y|x,d,m)\pi(m)dm$$

- Caveat: BDT encounters difficulties in infinite-dimensional parameter spaces (Diaconis & Freedman (1986)). So this BMA approach is not well suited to handle 'unstructured' model uncertainty.

Robust Control

- In Robust Control, agents do not adhere to the Savage axioms. They adhere to alternative axioms developed by Klibanoff, Marinacci, & Mukerji (2005) and Strzalecki (2011).
  1. They cannot form a unique prior.
  2. They do not reduce compound lotteries.
  3. They distinguish between model uncertainty and parameter uncertainty.
  4. They care about early resolution of uncertainty.
- All these behavioral modifications are controlled by a *single* free parameter, θ. This parameter is calibrated using detection error probabilities. As $\theta \to \infty$, Rational Expectations/Bayesian results are recovered.
- Robust Control features two models, not one. One is the agent's benchmark model. It is exogenously specified. The other is his worst-case model.
- Since the worst-case model depends on his own policy, the agent views himself as being immersed in a dynamic zero-sum game. A hypothetical 'evil agent' picks models to minimize his payoff.
- The agent is prudent, not paranoid. To avoid undue pessimism, the evil agent must pay a relative entropy cost.
- It is typically assumed the agent's doubts are all in his head. His benchmark model is correct. He just doesn't know it.

Robust Control: Mechanics

Benchmark Model

$$dx = f(x,c)dt + \sigma(x,c)dB$$

Relative Entropy

$$\mathcal{R}(q) = r\int_0^\infty e^{-rt}\left[\int \log\left(\frac{dq_t}{dq_t^0}\right)dq_t\right]dt$$

$$\int \log\left(\frac{dq_t}{dq_t^0}\right)dq_t = \frac{1}{2}\hat{E}\int_0^t |h_s|^2 ds$$

$$\hat{B}(t) = B(t) - \int_0^t h_s ds$$

Drift Distorted Models

$$dx = [f(x,c) + \sigma(x,c)h]dt + \sigma(x,c)d\hat{B}$$

Dynamic Zero-Sum Game

$$V(x) = \max_c \min_h \hat{E}_0 \int_0^\infty [u(x,c) + \frac{1}{2}\theta h^2]e^{-rt}$$

Worst-Case Model

$$h = -\frac{\sigma(c,x)}{\theta} V_x(x)$$

Ex. 1: Robust Portfolio Choice

- Let's now introduce model uncertainty into the Merton problem. Suppose the agent's benchmark model is:
$$\frac{dS}{S} = \mu dt + \sigma dB$$

- When discussing learning, we assumed the agent fully trusted this model. He just didn't know μ.

- Now suppose he fears more general forms of model misspecification (e.g., omitted variables, neglected nonlinearities, etc.)

- In response, he surrounds the benchmark model with a 'cloud' of alternative models, parameterized by h
$$\frac{dS}{S} = (\mu + \sigma h_t)dt + \sigma d\hat{B}$$
Notice that we have changed probability measures, from dB to $d\hat{B}$.

- To formulate robust saving/portfolio policies, he solves the following dynamic zero-sum game

$$\max_{c,\pi} \min_{h} \hat{E}_0 \int_0^\infty e^{-\delta t} \left[\frac{C_t^{1-\gamma}}{1-\gamma} + \frac{1}{2\psi} h_t^2 \right] dt$$

s.t. $$dW = [(r + \pi(\mu + \sigma h_t - r))W - C]dt + \pi \sigma W d\hat{B}$$

where for notational convenience, $\psi \equiv \theta^{-1}$.

- Applying Ito's lemma, the stationary HJB/Isaacs equation is

$$\delta V = \max_{c,\pi} \min_{h} \left\{ \frac{C^{1-\gamma}}{1-\gamma} + \frac{1}{2\psi} h^2 + [(r + \pi(\mu + \sigma h - r))W - C]V_w + \frac{1}{2}\pi^2 \sigma^2 W^2 \cdot V_{ww} \right\}$$

- The FOCs are:

$$\begin{aligned} \psi^{-1} h + \pi \sigma W V_w &= 0 \qquad (h) \\ C^{-\gamma} - V_w &= 0 \qquad (C) \\ (\mu + \sigma h - r)W \cdot V_w + \pi \sigma^2 W^2 V_{ww} &= 0 \qquad (\pi) \end{aligned}$$

Notice that the agent's plans depend on the actions of the 'evil agent' (and vice versa).

- Solving for h we get

$$h = -\psi\pi\sigma W \cdot V_w$$

This tells us a lot about robustness:

1. Notice that as long as $V_w > 0$ and $\pi > 0$, robustness takes the form of pessimism, i.e., the drift distortion is negative, $h < 0$.
2. The agent worries more when he is more 'exposed', i.e., when π is large.
3. The agent worries more when volatility, σ, is higher. Greater volatility makes it harder to rule out alternative models.
4. The agent's doubts depend on his wealth level. To a 1st-order approx., they decline with wealth when $\gamma > 1$

- Subbing this into the agent's FOC for π we get the robust portfolio policy

$$\pi = \left[\frac{-V_w}{W(V_{ww} - \psi V_w^2)}\right] \frac{\mu - r}{\sigma^2}$$

Notice that when $\psi = 0$ we get the usual Merton result. However, when $\psi > 0$ there is a sense in which the agent is 'more risk averse', since $V_{ww} - \psi V_w^2$ is more negative.

- However, unless $\gamma = 1$, so that $V \sim \log(W)$, the agent's problem is now nonhomothetic. Portfolio and savings policies will depend nonlinearly on the level of wealth.

- Nonhomotheticity is a general property of robust decision rules. It creates challenges for aggregation. It is also the subject of much recent research, e.g., Anderson (2005), Bhandari (2018), Borovicka (2018).

- Maenhout (2004) proposed a ‘trick’ to preserve homotheticity. He suggested scaling the robustness parameter by the inverse of the value function,

$$\hat{\psi} = \frac{\psi}{(1-\gamma)V}$$

When $\gamma > 1$, this implies agents become more ambiguity averse as their wealth increases. This keeps their doubts alive. [The reverse is true when $\gamma < 1$]. The decision-theoretic foundations of this trick remain unclear.

- Using this scaling trick, one can readily verify the following solution for the robust Merton problem

$$V(W) = A\frac{1}{1-\gamma}W^{1-\gamma} \qquad A = \gamma\left[\delta - (1-\gamma)r - \frac{1-\gamma}{2(\gamma+\hat{\psi})}\left(\frac{\mu - r}{\sigma}\right)^2\right]^{-1}$$

$$\pi = \frac{1}{\gamma+\hat{\psi}}\left(\frac{\mu - r}{\sigma^2}\right)$$

An Observational Equivalence

- Comparing this solution to our previous solution using Duffie-Epstein SDU preferences, we obtain the following intriguing observational equivalence result:

  Proposition: *An agent with time-additive CRRA preferences and a homothetic preference for robustness, $\hat{\psi}$, is observationally equivalent to a Stochastic Differential Utility agent with $\boldsymbol{EIS} = \gamma^{-1}$ and relative risk aversion $\gamma + \hat{\psi}$*

- Skiadas (2003) and Hansen, Sargent, Turmuhambetova, & Williams (*JET*, 2006) show that this observational equivalence extends beyond the realm of the Merton problem. It becomes more apparent if we use the unnormalized SDU aggregator. The $\hat{\psi}\boldsymbol{\sigma}^{2}(\boldsymbol{x})\boldsymbol{V}_{\boldsymbol{x}}^{2}$ term that appears in the robust HJB equation turns out to be the SDU variance multiplier.

- Observational equivalence is often regarded as a problem. For us it is an opportunity. This is because γ and $\hat{\psi}$ have different interpretations.

- γ reflects aversion to gambles with known probabilities, and can therefore be calibrated to match the low values often ascribed to these situations. In contrast, $\hat{\psi}$ reflects aversion to model uncertainty. It is an aspect of the environment, not preferences. Hence, it is context specific. Hansen and Sargent calibrate it using detection error probabilities.

Detection Error Probabilities

- In robust control, agents are prudent, not paranoid. They only worry about empirically plausible alternatives.

- Det. Error Prob
$$\frac{1}{2}P(H_1|H_0) + \frac{1}{2}P(H_0|H_1)$$

- Like a significance level, except null and alternative are symmetric.

- If DEP is small, then worst case model could be easily rejected based on historical data.

- *Example*: Two i.i.d. Gaussian processes with common variance
$$DEP \leq \frac{1}{2}\exp\{-T\rho\} \qquad \rho = \frac{(\mu_1 - \mu_0)^2}{8\sigma^2}$$

- DEP will be small when:
  1. Sample size, T, is large.
  2. Difference in means is large.
  3. Variance is small.

Some Results

- Remember, the Merton model is a partial equilibrium model. Prices are exogenous. However, we can gauge the impact of model uncertainty by looking at its effects on risky asset demand.

- To do this, let's calibrate parameters to conventional values, and then vary $\hat{\psi}$. The parameter values are: $\gamma = 5$, $\sigma = .16$, and $(\mu - r) = .06$ (corresponding to an annual time unit).

  The following table displays the resulting risky portfolio share as a function of $\hat{\psi}$

ψ	π	EP^Q	DEP (Chernoff)	DEP(Normal)
0	.469	.06	.500	.500
1	.391	.05	.476	.377
5	.234	.03	.322	.174

- The column EP^Q is the perceived equity premium. The last 2 columns display detection error probabilities. DEP(Chernoff) is the Chernoff bound. The final column uses the exact DEP, which is available due to the i.i.d./normal environment.

- Evidently, an empirically plausible fear of model misspecification can reduce the demand for risky assets by more than 50%! Current research is using ambiguity and model uncertainty to explain the lack of participation in financial markets.

Ex. 2: Equilibrium Prices

- Following the usual strategy, let's now take consumption as exogenous and compute equilibrium prices. Not surprisingly, the equilibrium response of a reduced demand for the risky asset is an increase in the risk premium.

- With homothetic robustness, the economy remains i.i.d., so to compute market-clearing prices we can just impose the market-clearing condition $\pi = 1$ in the Merton soln.

$$\begin{aligned} r &= \delta + \gamma\mu_c - \frac{1}{2}(1+\gamma)(\gamma+\hat{\psi})\sigma_c^2 \\ \mu_S - r &= (\gamma+\hat{\psi})\sigma_c^2 \end{aligned}$$

- Note that we can keep γ low to match the risk-free rate, while increasing $\hat{\psi}$ as needed in order to match the equity premium. The only real question is whether the implied DEP is plausible.

- Barillas, Hansen & Sargent (*JET*, 2009) examine US quarterly data for the period 1948-2007. They assume $IES = 1$ and suppose the benchmark model for consumption growth is i.i.d. They find that the SDF nearly hits the HJ bound for $DEP = .05$. It gets half-way to the HJ bound for $DEP \in (.15 - .20)$.
- Which is more plausible? A coefficient of relative risk aversion of 50, or a DEP of .05?

Recent Work

This lecture has reviewed what can be considered '1st-generation' models of robust control and model uncertainty. Here is quick overview of more recent work:

1. Like the rest of economics, recent work focuses more on heterogeneity and aggregation:
   - Bhandari (2014), *Doubts, Asymmetries, and Insurance* – 2 agents with distinct models. Time-varying Pareto weights. Stable wealth dist if $IES > 1$.
   - Miao & Rivera (Ecma, 2016), *Robust Contracts* – Principal agent problem. Principal has doubts about the cash flow process.
   - Kasa & Lei (JME, 2018), *Risk, Uncertainty, & the Dynamics of Inequality* – OLG model with nonhomothetic robustness. Wealthy agents are less ambiguity averse, so invest more in higher yielding assets. This amplifies and accelerates wealth inequality.
   - Cho & Kasa (2018), *Doubts, Inequality, & Bubbles* – Two agents trade an indivisible asset, and have doubts about underlying cash flows. Doubts depend on wealth, which depends on asset ownership. Heterogeneous beliefs create a resale option value as in Scheinkman & Xiong (2003).

2. Disentangling intertemporal substitution, risk aversion, and ambiguity aversion.
   - The above noted observational equivalence has sparked a recent literature that attempts to distinguish risk aversion from both ambiguity aversion and intertemporal substitution. The key to separating risk aversion from ambiguity aversion is to introduce hidden state variables, which the agent attempts to learn about. Ambiguity is defined by distortions of the agent's estimates of the hidden states. Hidden states can be used to represent alternative models.
   - Distorted beliefs about hidden states can be interpreted from the perspective of the KMM (2005) model of smooth ambiguity aversion. In the KMM model an agent prefers act f to act g if and only if

$$\mathbb{E}_{\mu}\phi(\mathbb{E}_{\pi}u \circ f) \geq \mathbb{E}_{\mu}\phi(\mathbb{E}_{\pi}u \circ g)$$

     where $\mathbb{E}$ is the expectation operator, π is a probability measure over outcomes conditional on a model, and μ is a probability measure over models. Ambiguity aversion is characterized by the properties of the ϕ function, while risk aversion is characterized by the properties of the u function. If ϕ is concave, the agent is ambiguity averse.

- KMM refer to $\mathbb{E}_{\pi}$ as 'first-order beliefs', while $\mathbb{E}_{\mu}$ is referred to as 'second-order beliefs'. Note that when ϕ is nonlinear, the implicit compound lottery defined by selecting a model with unknown parameters cannot be reduced to a single lottery over a 'hypermodel', as in Bayesian decision theory, so the distinction between models and parameters becomes important. Also note that from the perspective of smooth ambiguity aversion, evil agents and entropy penalized drift distortions are just a device used to produce a particular distortion in second-order beliefs about continuation values, i.e., where $\phi(V) \approx -\exp(-\varepsilon V)$.
- The KMM model is static. Hayashi & Miao (2011) propose a model of generalized smooth ambiguity aversion by combining KMM with an Epstein-Zin aggregator. Unfortunately, this model does not extend to continuous-time with Brownian information structures. Intuitively, first-order uncertainty (risk) is $O(dt)$, whereas second-order uncertainty (ambiguity) is $O(dt^2)$, and so it evaporates in the continuous-time limit. In response, Hansen & Sargent (*JET*, 2011) and Hansen & Miao (2018) propose a trick to retain ambiguity aversion, even as the sampling interval shrinks to zero. In particular, they show that if the robustness/ambiguity-aversion parameter is scaled by the sampling interval, then ambiguity aversion will persist in the limit. Intuitively, even though second-order uncertainty becomes smaller and smaller as the sampling interval shrinks, because the agent effectively becomes more ambiguity averse at the same time, ambiguity continues to matter.

3 Learning

- Learning and robustness are alternative responses to model uncertainty. With learning you try to reduce uncertainty. With robustness you simply cope with it.
- A natural question is which is the 'better' approach. Hansen & Sargent (2011, *Wanting Robustness in Macroeconomics*) argue that it depends on the dimensionality of the problem. You can learn your way out of small dimensional uncertainty, but not infinite-dimensional uncertainty. (Remember Savage's warning about 'small worlds'!)
- Hansen & Sargent (*JET*, 2007) develop a framework that combines learning and robust control. An agent entertains a small number of competing models, each of which is potentially misspecified. Over time he (robustly) learns which is the best of the potentially misspecified models. In their *Fragile Beliefs* paper, HS apply this framework to the Bansal-Yaron long-run risks model. The agent does robust BMA over a model with i.i.d. consumption growth, and an alternative with persistent growth fluctuations. They show that weight on the LRR model is countercyclical, the agent assigns higher probability to the LRR model during bad times. Essentially, the agent thinks bad times will be persistent, while good times will be transitory. This generates a large countercyclical risk premium even though actual consumption growth is i.i.d.!