

A Bayesian Model of Knightian Uncertainty

Nabil I. Al-Najjar *and* Jonathan Weinstein

D-Tea, Paris May 2013

Presentation also draws from ideas in:

- *The Ambiguity Aversion Literature: A Critical Assessment (with Jonathan Weinstein)*
- *A Bayesian Approach to Precautionary Policies*
- *Uncertainty and Disagreement in Equilibrium Models (with Eran Shmaya)*

Plan of the Talk

► *Skip slide*

- 1 Introduction
- 2 Precautionary saving under parameter uncertainty
- 3 Testing as criterion for objective/subjective probability
- 4 Rational expectation econometrics and the “problem” of Knightian uncertainty

Two types of probabilities

► Skip slide

Knight, 1921 argued there is a fundamental difference between

“measurable risk and an unmeasurable uncertainty”

Intuition: A statements like

- *The probability that it rains tomorrow is xx %*

seems more “objective,” or less likely to provoke disagreement than statements about the probability of

- *US debt-to-GDP ratio drops below 60% in 2030*
- *The European monetary union unravel within 5 years*

Many people “feel” that these probability judgments are fundamentally different in nature

Keynes on Uncertainty

► Skip slide

Keynes, 1937:

“The sense in which I am using the term [uncertainty] is that in which the prospect of a European war is uncertain, or the price of copper and the rate of interest twenty years hence, or the obsolescence of a new invention, or the position of private wealth-owners in the social system in 1970.”

About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know.”



Is this just semantics?

► *Skip slide*

Perhaps, ... but it is semantics that increasingly resonates:

- Model-uncertainty
- Disagreement and heterogenous beliefs
- Precautionary policies in the context of “deep uncertainty”

Commonly held point of view:

- The Bayesian framework cannot capture these phenomena

Routledge and Zin on Knightian Uncertainty

► Skip slide

“Savage rationality implies that [...] ‘model uncertainty,’ is indistinguishable from the risk inherent in the asset’s stochastic process.

The Savage independence axiom implies that one can simply collapse the probability weighting across possible models (“uncertainty”) with the probabilities for payoffs (“risk”) to represent behavior with a single probability measure for states.”

A Bayesian Framework for Knightian Uncertainty

- Bayesian? *You gotta be kidding..*
- *“Why go back to tired old ideas?”*
- Let's take a moment to revisit some intellectual history...

Knight on Knightian uncertainty

http://bit.ly/Knight_on_Knightian_Uncertainty

Keynes on Knightian Uncertainty

► Skip slide

“when an individual instance only is at issue, there is no difference for conduct between a measurable risk and an unmeasurable uncertainty.

The individual [in this case] throws his estimate of the value of an opinion into the probability form and ‘feels’ toward it as toward any other probability situation.”

Knight, 1921



Keynes on Knightian Uncertainty

► Skip slide

“Nevertheless, the necessity for action compels us to behave exactly as if we had behind us [...] prospective advantages and disadvantages, each multiplied by its appropriate probability, waiting to be summed.”



Two Approaches

► *Skip slide*

Common intuitions are:

- People dislike uncertainty
- People tend to display excessively precautionary behavior
 - “excessive” relative to what?
- Two approaches:
 - *Agents dislike uncertainty as a matter of taste*
 - *Agents dislike uncertainty because it makes planning harder*

Income and consumption

► Skip slide

Permanent income model a la Cogley - Sargent (IER, 2008)

- **Consumption:** Infinite horizon and time-separable utility:

$$U(c_1, \dots) = \sum_{i=1}^{\infty} \delta^i u(c_i)$$

- **Income** in period t is f_t
 - Assumes finite possible values
 - f_t is distributed i.i.d. P^θ

Consumption plans

► Skip slide

- A **consumption plan**

$$c_t : H_t \rightarrow \mathbb{R}$$

- H_t history up to time t

- **Intertemporal budget constraint**

$$w_{t+1} = (1 + r)(w_t + f_t - c_t).$$

with $\delta = 1/1 + r$

Uncertainty

► Skip slide

- If the agent does not know the true i.i.d. parameter θ , he treats it as random with subjective belief μ
- μ induces a distribution P^μ on infinite sequences

Risk = known θ :

Uncertainty = unknown θ :

- More generally, ergodic vs. non-ergodic probabilities
- *How is uncertainty different from ordinary incomplete information?*
 - Answer: testing and estimation

Direct utility blurs the distinction

► Skip slide

Uncertainty on *consumption streams*:

- Implicitly rules out intertemporal smoothing
- $f_i = c_i$ in all states

Uncertainty reduces to risk:

$$\begin{aligned}
 E_{\mu} E_{\theta} \sum_{i=1}^{\infty} \delta^i u(c_i) &= \sum_{i=1}^{\infty} \delta^i E_{\mu} E_{\theta} u(c_i) \\
 &= E_{\bar{\theta}^{\mu}} \sum_{i=1}^{\infty} \delta^i u(c_i).
 \end{aligned}$$

$\bar{\theta}^{\mu}$ is the “average parameter;” *i.e.*, i.i.d. distribution with the same marginal as P^{μ}

Preliminaries

► Skip slide

Assumption: CARA period utility: $u(x) = -e^{-ax}$

The Euler equation for this problem is

$$u'(c_t^*) = E_t[u'(c_{t+1}^*)].$$

For the exponential utility, marginal utility is a constant multiple of the utility level, so for every k

$$u(c_t^*) = E_t[u(c_{t+k}^*)]$$

Thus the value function is the same as the instantaneous utility along an optimized plan:

$$V(w_t, \mu_t) = (1 - \delta) E_{\mu_t} \left[\sum_{l=t}^{\infty} \delta^l u(c_l^*) \right] = u(c_t^*).$$

Consumption under risk

► Skip slide

Proposition: Consumption when the true θ is known follows the autoregression:

$$c_{t+1} = c_t + \Gamma(\theta) + r(f_t - E[f|\theta])$$

where

$$\Gamma(\theta) \equiv E(rf|\theta) - CE(rf|\theta) \geq 0.$$

is the risk premium of the random variable rf , representing drift in consumption, and $r(f_t - E[f|\theta])$ is mean-zero error term.

Beliefs play no role; standard dynamic programming.

Caballero (1990); Kimball (several papers); Carroll (several papers); Cogley and Sargent (2008)

Uncertainty

► Skip slide

Under CARA, adding wealth simply results in additional interest income being consumed each period:

Lemma:

$$c^*(\mu, w) = \frac{r}{1+r} w + c^*(\mu, 0).$$

Notation:

$$c^*(\mu) \equiv c^*(\mu, 0) \quad \text{and} \quad V(\mu) \equiv V(\mu, 0).$$

Cost of uncertainty

► Skip slide

Proposition: For any w , V and c^* are strictly convex in μ . That is, letting μ_1 and μ_2 be any beliefs which result in distinct consumption plans,

- 1 $V(\lambda\mu_1 + (1 - \lambda)\mu_2, w) < \lambda V(\mu_1, w) + (1 - \lambda)V(\mu_2, w);$
- 2 $c^*(\lambda\mu_1 + (1 - \lambda)\mu_2, w) < \lambda c^*(\mu_1, w) + (1 - \lambda)c^*(\mu_2, w).$

Furthermore, for any μ with finite support, letting $\bar{\theta}^\mu$ be the average of θ under μ , if $r < 1$ then

$$V(\mu) \leq V(\bar{\theta}^\mu)$$

with strict inequality whenever μ is not a point mass, implying also

$$c^*(\mu) \leq c^*(\bar{\theta}^\mu).$$

Evolution of consumption under uncertainty

► Skip slide

Proposition: Under uncertainty, consumption evolves according to:

$$c_{t+1} = c_t + \Gamma(\mu) + r [f_t - E_{\mu_t}[f_t]] \\ + [E_{\mu_t}[c^*(\mu_{t+1})] - c^*(\mu_t)] + [c^*(\mu_{t+1}) - E_{\mu_t}[c^*(\mu_{t+1})]].$$

Intuition: uncertainty messes up intertemporal planning; agent prefers early resolution of uncertainty (not as a matter of taste, but for instrumental reasons). Agent perceives income shocks to have a permanent component (because beliefs are a martingale) and saves way more initially....

The first line in closely resembles behavior under risk. The second line includes two new terms due to the resolution of uncertainty and therefore absent under risk:

- Upward drift in consumption is due to the resolution of uncertainty which implies less precautionary savings. Bayesian updating implies $\mu_t = E_t[\mu_{t+1}]$, and by Jensen's inequality and the convexity of c^* , $E_t[c^*(\mu_{t+1})] \geq c^*(\mu_t)$
- New source of randomness due to the persistent effect of the random resolution of uncertainty.

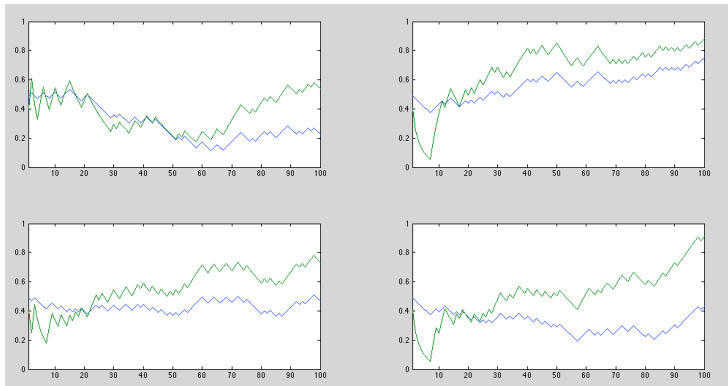
Simulation

► *Skip slide*

Even this simple model is intractable analytically.

As illustration of the solution, simulate the model assuming


- 1 Riskless rate: $r=0.04$
- 2 Risk aversion: $a=2$
- 3 Income: $f \in \{0, 1\}$ i.i.d.
- 4 Agent's subjective belief μ is uniform over θ 's



Four samples drawn under assuming true parameter: $\bar{\theta} = 0.50$

In each panel,

- 1 Same income realization, drawn from $\bar{\theta} = 0.50$
- 2 Blue line is consumption if the agent knew $\bar{\theta}$
- 3 Green line is consumption under the subjective belief μ



Average consumption over 10,000 samples

- Blue line is consumption if the agent knew $\bar{\theta}$
- Green line is consumption if the agent's subjective belief μ

- 1 Introduction
- 2 Precautionary Saving
- 3 Objectivity and Testing**
- 4 Rational Expectation Econometrics

The Borel Criterion

► Skip slide

How is subjective uncertainty different from objective risk?

Our answer is based on the Borel Criterion (Dawid (2004)):

- Uncertainty has observable implication on behavior
- but it is subjective in the sense that it cannot be tested against objective data

(“About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know.”)

Statistical tests

► Skip slide

Consider the simplest example of uncertainty: only two parameters $\theta_1 \neq \theta_2$

$\hat{\Delta}$ is the set of beliefs μ with $\mu(\theta_i) > 0, i = 1, 2$.

Abstractly, a statistical test is a function

$$T : \{\theta_1, \theta_2\} \times \mathcal{S}^\infty \rightarrow \{0, 1\}$$

- $T(\theta_m, \mathbf{s}^\infty) = 1$: the data \mathbf{s}^∞ confirms the hypothesis that the true process is θ_m
- $T(\theta_m, \mathbf{s}^\infty) = 0$: the data \mathbf{s}^∞ is inconsistent with θ_m

Asymptotic testing

► Skip slide

To focus on what is *in principle* testable, consider asymptotic tests that make use of the entire infinite sequence of data

With unlimited data, it is natural to consider tests that are:

- 1 *Unprejudiced*: T has Zero Type I error on parameters

$$P^{\theta_m} \{ T(\theta_m, \mathbf{s}^\infty) = 1 \} = 1$$

- 2 *Powerful*: T has Zero Type II error on parameters

$$P^{\theta_{m'}} \{ T(\theta_{m'}, \mathbf{s}^\infty) = 1 \} = 0$$

Testing uncertain beliefs

► Skip slide

Extend T to $\hat{\Delta}$, the set of all beliefs with (strict) support $\{\theta_1, \theta_2\}$

$$\hat{T} : \hat{\Delta} \times \mathcal{S}^\infty \rightarrow \{0, 1\}$$

such that \hat{T} has zero Type I error on $\hat{\Delta}$

Proposition: Let μ, μ' be any pair of beliefs in $\hat{\Delta}$ and \hat{T} any test with the above properties. Then

$$P^\mu \{ \hat{T}(\mu', \mathbf{s}^\infty) = 1 \} = P^{\mu'} \{ \hat{T}(\mu, \mathbf{s}^\infty) = 1 \} = 1$$

An agent who believes the data is generated by μ_j also believes that no zero-Type-I-error-test can disprove the wrong belief μ_j .

How should Rational Expectations be defined in this case?

Proof

► Skip slide

Zero Type I error on uncertain beliefs means:

$$1 = P^{\mu'} \{ \hat{T}(\mu', \mathbf{s}^\infty) = 1 \}$$

$$= P^{\theta_1} \{ \hat{T}(\mu', \mathbf{s}^\infty) = 1 \} \mu'(\theta_1) + P^{\theta_2} \{ \hat{T}(\mu', \mathbf{s}^\infty) = 1 \} \mu'(\theta_2).$$

$$\implies$$

$$P^{\theta_1} \{ \hat{T}(\mu', \mathbf{s}^\infty) = 1 \} = P^{\theta_2} \{ \hat{T}(\mu', \mathbf{s}^\infty) = 1 \} = 1$$

$$\implies$$

$$\begin{aligned} P^\mu \{ \hat{T}(\mu', \mathbf{s}^\infty) = 1 \} &= P^{\theta_1} \{ \hat{T}(\mu', \mathbf{s}^\infty) = 1 \} \mu(\theta_1) + P^{\theta_2} \{ \hat{T}(\mu', \mathbf{s}^\infty) = 1 \} \mu(\theta_2) \\ &= 1 \end{aligned}$$

The other equality is proved similarly.

Finite tests

► Skip slide

If the test has access to finite data, a similar (weaker) conclusion holds.

- T can condition on only n observations
- Require n to be large enough to have ϵ -Type I and II errors
- μ puts 0.5 on each of θ_1, θ_2
- ϵ Type I error means:

$$P^{\theta_1}\{T(\mu', \mathbf{s}^\infty) = 1\}\mu'(\theta_1) + P^{\theta_2}\{T(\mu', \mathbf{s}^\infty) = 1\}\mu'(\theta_2) > 1 - \epsilon$$

$$P^{\theta_1}\{T(\mu', \mathbf{s}^\infty) = 1\} + P^{\theta_2}\{T(\mu', \mathbf{s}^\infty) = 1\} > 2 - 2\epsilon$$

$$P^{\theta_m}\{T(\mu', \mathbf{s}^\infty) = 1\} > 1 - 2\epsilon$$

Tests of uncertain beliefs have little power

► Skip slide

- Take any μ' that puts mass α on P^{θ_1} and $1 - \alpha$ on P^{θ_2}
- Then an agent with beliefs μ' will also believe that μ must pass the test:

$$\begin{aligned}\mu'\{T(\mu, \omega) = 0\} &= \alpha P^{\theta_1}\{T(\mu, \omega) = 1\} + (1 - \alpha) P^{\theta_2}\{T(\mu, \omega) = 1\} \\ &> 1 - 2\epsilon\end{aligned}$$

- Tests have little power, regardless of the nature of the test or the amount of data

- 1 Introduction
- 2 Precautionary Saving
- 3 Objectivity and Testing
- 4 Rational Expectation Econometrics**

So far so good...

► *Skip slide*

But why should uncertainty about θ deserve special treatment, different from plain old imperfect information models?

Uncertainty is difficult to reconcile with the Rational Expectations paradigm

- 1 Seen this with testing: requiring RE beliefs to be testable excludes uncertainty
- 2 Uncertainty is problematic to Rational Expectations Econometrics

Uncertainty and Rational Expectations Econometrics

REE is the backbone of most intertemporal empirical models

- Empirical IO
- Macro
- Asset pricing

Powerful methodology to remove expectations as a free variable via the identifying assumption:

“Beliefs coincide with observed frequencies.”

REE in the consumption-saving model

► Skip slide

Simple model:

- 1 $f \sim i.i.d.(\theta)$

- 2

$$c_{t+1} = c_t + \Gamma(\mu) + r [f_t - E_{\mu_t}[f_t]] \\ + [E_{\mu_t}[c^*(\mu_{t+1})] - c^*(\mu_t)] + [c^*(\mu_{t+1}) - E_{\mu_t}[c^*(\mu_{t+1})]]$$

An estimator is a function: Data $\mapsto (\hat{\theta}, \hat{a}, \dots)$

However, the agent's posterior μ_t is *unobservable*

Even if observed, one would like to connect beliefs with reality, to avoid “*anything goes!*”

Something has to be done

RE identifying assumption:

“beliefs=econometrician’s estimate of $\bar{\theta}$.” $\mu = \hat{\theta}$

Imagine applying a GMM estimator to the model:

$$\begin{aligned}c_{t+1} = & c_t + \Gamma(\hat{\theta}) + r [f_t - E_{\hat{\theta}}[f_t]] \\ & + [E_{\hat{\theta}}[c^*(\hat{\theta})] - c^*(\hat{\theta})] + [c^*(\hat{\theta}) - E_{\hat{\theta}}[c^*(\hat{\theta})]].\end{aligned}$$

With simple algebra, we are back to evolution under risk

$$c_{t+1} = c_t + \Gamma(\hat{\theta}) + r [f_t - E_{\hat{\theta}}[f_t]]$$

REE in the precautionary saving model

► Skip slide

At the end there is only one true $\bar{\theta}$

Rational Expectations (my interpretation of Cogley-Sargent):

- There is large amount of (unmodeled) past data, observed by agent but not the econometrician
- Agents believe the future will look like the past (*i.e.*, they believe they know the ergodic state)
- The econometrician's estimate $\hat{\theta}$ is then a good estimate of agents' beliefs

This rules out uncertainty about parameters as possible motive for choice and may result in a mis-specified model

- Weitzman (2007)'s critique of asset pricing models

- Under risk: REE works amazingly well
- But applying REE when agents have uncertain beliefs may give misleading impressions
 - 1 Agents are irrational
 - 2 ...have excessive risk aversion
 - 3 ...do not have a well-defined prior
 - 4 ...etc.
- But doing time series econometrics under uncertainty is *very hard*.

Conclusion

► *Skip slide*

Decision theoretic approach: agents dislike uncertainty, *as a matter of taste*

In a Bayesian framework uncertainty corresponds to properties of the *indirect* utility, reflecting constraints of information, lending and borrowing imperfections... etc

But this raises significant conceptual and econometric modeling challenges