

A Bayesian Model of Knightian Uncertainty

Nabil I. Al-Najjar* *and* Jonathan Weinstein†

First draft: December 2011

This version: September 2012

Abstract

A long tradition suggests a fundamental distinction between situations of risk, where true objective probabilities are known, and unmeasurable uncertainties where no such probabilities are given. This distinction can be captured in a Bayesian model where uncertainty is represented by the agent's subjective belief over the parameter governing future income streams. Whether uncertainty reduces to ordinary risk depends on the agent's ability to smooth consumption. Uncertainty can have a major behavioral and economic impact, including precautionary behavior that may appear overly conservative to an outside observer. We argue that one of the main characteristics of uncertain beliefs is that they are not empirical, in the sense that they cannot be objectively tested to determine whether they are right or wrong. This can confound empirical methods that assume rational expectations.

* Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, Evanston IL 60208. [Al-Najjar's Research Page](#).

† Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, Evanston IL 60208. [Weinstein's Research Page](#).

Contents

1	Introduction	1
2	Risk and Uncertainty	5
2.1	Model and Notation	5
2.2	Risk, Uncertainty, and the Value Function	6
3	Uncertainty and Precautionary Behavior	10
3.1	Model	10
3.2	Pure Risk	11
3.3	Uncertainty	13
4	Empirical Implications of Uncertainty	15
4.1	Uncertain Beliefs are Untestable	15
4.2	Inference under Uncertainty	17
4.3	Discussion and Related Literature	21
A	Appendix	24

1 Introduction

Knight (1921)'s idea of a fundamental difference between “measurable risk and an unmeasurable uncertainty” has generated both interest and controversy. Standard models in economics assume that agents use probabilities to quantify all uncertainties regardless of their source or nature. No distinction is drawn between actuarial and strategic risks, or between risks associated with repetitive vs. singular events. Yet there is a compelling intuition that some probability judgments are less ‘obvious’ or ‘objective’ than others. Betting on events like the unraveling of the European monetary union or global warming seems qualitatively different from betting on the outcome of a coin toss or whether it rains tomorrow. The first type of event represents, in Knight’s words, unmeasurable uncertainties that should be treated differently from measurable risks.

This paper suggests that a Bayesian framework can capture the distinction between situations with known probabilities, or “risk,” and Knightian uncertainty where objective probabilities are unknown.¹ We illustrate the distinction in a consumption-saving model where uncertainty is represented by the agent’s subjective belief about the parameter governing future income. We relate the impact of uncertainty to the agent’s ability to intertemporally smooth consumption and show that this leads to precautionary behavior that may appear overly conservative to an outside observer. We also point to the potential tension between subjective uncertainty and empirical methods that use rational expectations for econometric identification.

Consider an agent with time-separable utility over (finite or infinite) consumption streams. The agent receives an i.i.d. income stream with unknown parameter θ .² Uncertainty is lack of knowledge represented by a prior belief μ over θ . Uncertainty does not necessarily lead to measurable consequences

¹ Keynes (1937) characterized uncertain beliefs as follows: “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.”

² The i.i.d. assumption simplifies the analysis, but is not essential to our point. We discuss this point in Section 4.

on behavior. For example, if uncertainty is defined on consumption streams directly, then discounted expected utility is unaffected by replacing the uncertain belief μ by a parameter with the same marginal as μ . In this case, uncertainty reduces to ordinary risk. The role of uncertainty is manifested in the indirect utility the agent derives from income streams when he has (perhaps limited) freedom to save or borrow.

We first illustrate this point in a simple setting where consumption occurs after an initial phase of payoff accumulation. For example, a retirement portfolio generates dividends each period, but the agent cares only about its value at the time of retirement. Another example is a start-up that accumulates gains and losses over a period of time, but whose value is realized only when the entrepreneur sells the firm. In these examples, separating consumption from payoff accumulation simplifies the calculation of indirect utility, making the impact of uncertainty transparent and striking.

We then turn to a richer consumption-saving problem where consumption and wealth can change as uncertainty about income resolves. Using exponential utility for tractability, we derive the evolution of consumption and show that uncertainty results in precautionary behavior that may appear overly conservative to an outside observer. Consumption is more volatile under uncertainty because the agent perceives short-run income variations to be potentially informative about his long-run income prospects. Under risk, by contrast, all income realizations are viewed as transitory.

A Bayesian framework represents any lack of knowledge in terms of probabilities. What then justifies treating uncertainty differently from other situations of imperfect information? Our main point is that the key difference between risk and uncertainty is that uncertain beliefs are not *empirical*. Section 4 introduces two arguments to support this point.

First, we formalize the intuition that the probabilities of uncertain events are subjective opinions about which, in Keynes' words, "there is no scientific basis to form any calculable probability whatever." We capture this intuition using statistical tests that compare agents' subjective beliefs with the actual sequence of realized outcomes. Consider asymptotic tests for simplicity. A natural property to require in such tests is to be free of Type I error: if the agent knows the true probabilities, he must pass the test almost surely.

Proposition 4.1 says that an agent who is uncertain about the true parameter, and who has subjective belief μ , must also believe that there is an alternative belief $\mu' \neq \mu$ such that no Type I error free test could reject μ' regardless of the amount of data used. Bayesian agents assign probabilities to *all* events, whether risky or uncertain. What distinguishes beliefs about uncertain events is that they cannot be objectively tested to determine whether they are right or wrong. By contrast, it is easy to test beliefs under risk by comparing them with observed frequencies.

The second sense in which uncertain beliefs are not empirical concerns the difficulty of estimating their impact using standard econometric methods. Beliefs influence decisions regardless of whether they reflect risk or uncertainty. But since beliefs are not directly observable, econometric identification assumptions are needed to recover them from data. A standard assumption is rational expectations which identifies beliefs with observed empirical frequencies. While this assumption offers considerable advantages, it also rules out subjective model-uncertainty as a factor in decisions, a point made by Weitzman (2007) among others. In Section 4 we suggest that the perceived failure of equilibrium models to capture Knightian uncertainty may have more to do with the use of rational expectations in their econometric estimation than with the Bayesian rational choice paradigm.

There is a growing interest in the economic role of Knightian uncertainty.³ One motivation is the discrepancy between observed behavioral patterns (*e.g.*, in asset prices) and the predictions of models where agents are assumed to know the true data generating process. Introducing uncertainty about fundamentals is a natural way to bring models closer to reality. Pástor and Veronesi (2009) survey asset pricing anomalies that could be explained with the introduction of uncertainty about fundamental parameters. They conclude that “[m]any facts that appear baffling at first sight seem less puzzling once we recognize that parameters are uncertain and subject to learning.” Hansen and Sargent (2001) discuss the importance of model misspecification and parameter-uncertainty in macroeconomic modeling. Connections to these works are discussed in greater detail in Section 4.

³ The terms parameter-uncertainty, model-uncertainty, or model mis-specification are often used instead of what we simply call “uncertainty.”

Two seminal papers, Gilboa and Schmeidler (1989) and Bewley (1986), formalize the concept of Knightian uncertainty as lack of full Bayesian belief. In Bewley (1986), uncertainty is modeled as an incomplete ranking over acts. Gilboa and Schmeidler (1989)'s ambiguity averse agents use a maxmin criterion with respect to a set of priors to incorporate caution in their decisions. Both approaches focus on the typical pattern of choices in static Ellsberg experiments as a key behavioral manifestation of uncertainty.

The present paper argues that a distinction between risk and Knightian uncertainty can be made within the Bayesian framework. This point of view follows a number of authors, including Halevy and Feltkamp (2005) and Weitzman (2007), who pursue Bayesian approaches to uncertainty. LeRoy and Singell (1987) suggest that such approach can be traced to Knight (1921)'s original work, noting that "Knight shared the modern view that agents can be assumed always to act as if they have subjective probabilities." Similarly, Keynes (1937) writes that, even in situations of uncertainty, "the necessity for action and for decision compels us [...] to behave exactly as we should if we had [...] a series of prospective advantages and disadvantages, each multiplied by its appropriate probability, waiting to be summed." Knight and Keynes, writing decades before modern subjective expected utility theory, seemed to believe that decision making under uncertainty is not necessarily in conflict with probabilistic reasoning.

We find it useful to distinguish Knightian uncertainty from ambiguity aversion. We take uncertainty to mean probabilities that cannot be objectively measured, an intuition we formalize in terms of statistical tests. Ambiguity aversion, on the other hand, refers to non-probabilistic beliefs, exemplified by the static Ellsberg choices. Although both lead to precautionary behavior, there are profound differences. In a Bayesian model, the implications of uncertainty appear in connection with intertemporal choice and the constraints on consumption smoothing. In static settings, such as Ellsberg's choices, risk and uncertainty are indistinguishable. This consistent with Knight (1921)'s view: "when an individual instance only is at issue, there is no difference for conduct between a measurable risk and an unmeasurable uncertainty. The individual [...] throws his estimate of the value of an opinion into the probability form of ' a successes in b trials' [...] and 'feels'

toward it as toward any other probability situation.” In modern Bayesian language, agents care only about the prizes they receive, not whether they were the result of risk rather than uncertainty.

The outline of the rest of the paper is as follows. Section 2 introduces the basic setup and defines the uncertainty premium. This section also introduces a very simple model where the value function can be easily computed and the impact of uncertainty is obvious. This simple model is related to Halevy and Feltkamp (2005), which we discuss in some detail. Section 3 introduces a more complete saving-consumption model and derives the stochastic laws of consumption under risk and uncertainty. Uncertainty leads to precautionary behavior and greater sensitivity to information. Section 4 discusses the empirical implications of uncertainty. We begin with a simple argument illustrating that uncertain beliefs are not testable, then discuss the potential tension between uncertainty and rational expectations econometrics. We discuss in detail the relationship of our work to Weitzman (2007) and Cogley and Sargent (2008).

2 Risk and Uncertainty

2.1 Model and Notation

We consider infinite horizon decision problems (finite horizon problems can be obtained as a special case). In each period, an outcome in a finite set $S = \{s^1, \dots, s^k\}$ is realized. The set of infinite sequences of outcomes is denoted S^∞ . We use subscripts to indicate time periods and superscripts to indicate outcomes. For example, s_i is the outcome at time i while s^j is the j th outcome.

An agent has a time-separable discounted utility for consumption streams:

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

where $\delta \in [0, 1]$ is a discount factor and $u : \mathcal{R} \rightarrow \mathcal{R}$ is a concave von Neumann-Morgenstern utility.

Let Θ be the set of all probability distributions $\theta = (\theta^1, \dots, \theta^k)$ on S . In this paper, a *parameter* is an i.i.d. distribution P^θ on S^∞ obtained by independently sampling from S according to θ . We will refer to either θ or P^θ as “parameter.”⁴ We do not consider more general parametric models, such as Markov processes, for tractability and expositional simplicity.

The agent’s uncertainty about θ is represented by a prior μ over Θ . Let P^μ represent the implied belief about infinite samples, defined by $P^\mu(B) \equiv \int_{\Theta} P^\theta(B) d\mu(\theta)$, for every event B .⁵ From the perspective of the agent, the parameter θ is a random variable with distribution μ . Expectations with respect to μ and θ are denoted E_μ, E_θ , respectively.

It will be useful to define the “average” parameter $\bar{\theta}^\mu$ by $\bar{\theta}^\mu(s^i) = E_\mu \theta(s^i)$. For example, if s^i is ‘Heads’ in a coin toss, then the space of parameters is $[0,1]$ and μ is a distribution on $[0,1]$. In this case, $E_\mu \theta^i$ is the expected value of the probability of Heads. Although P^μ and $P^{\bar{\theta}^\mu}$ share the same marginal on any single coordinate, $P^{\bar{\theta}^\mu}$ is always independent, while P^μ is independent only when μ concentrates all its mass on a single parameter.

2.2 Risk, Uncertainty, and the Value Function

A Bayesian agent who is uncertain about the ‘true’ parameter θ represents this lack of knowledge in terms of a prior μ . His expected utility on consumption streams is:

$$E_\mu E_\theta U(c_1, \dots).$$

It is easy to see that this equals $E_{\bar{\theta}^\mu} U(c_1, \dots)$, so the agent is indifferent between uncertainty about the parameter and certain knowledge of the average parameter.⁶

This makes a simple but important reference point: how confident an agent ‘feels’ about his knowledge of the true θ is irrelevant in ranking con-

⁴ All events in S^∞ are assumed to be Borel sets. The space of probability measures on S^∞ is itself endowed with the weak* topology and the Borel sigma-algebra of events. All functions used in the paper are assumed measurable.

⁵ There is a 1-1 correspondence between P^μ and μ , yet they are different objects: the former is a distribution on S^n , while μ is a distribution on Θ .

⁶ This follows from the linearity of probabilities and the time-separability of utility: $\int_{\Theta} \int_{S^n} \sum_{i=1}^{\infty} \delta^i u(c_i) dP^\theta d\mu = \sum_{i=1}^{\infty} \delta^i E_{\bar{\theta}^\mu} u(c_i) = E_{\bar{\theta}^\mu} U(c_1, \dots, c_n)$.

sumption streams. This apparent failure of the standard framework to capture uncertainty about the true probabilities may be behind the perception that Knightian uncertainty requires a departure from probability-based reasoning. For instance, consider an agent A who believes that his consumption stream will either be high forever or low forever with equal probability, vs. an agent B who believes that each day there is an independent 50-50 draw between high and low. Our standard utility function evaluates these as equal, because agent B does not derive any hedging benefit from the independent draws, but is forced to “starve” on bad days. It is more natural, though to consider the case that some intertemporal smoothing is possible.⁷

While a Bayesian framework cannot distinguish risk and uncertainty in *consumption streams*, this distinction is possible, indeed natural, in evaluating the indirect utility of *payoff streams*. To make this idea formal, define a *consumption plan* c as a sequence of functions:

$$c_i : \mathcal{R}^i \rightarrow \mathcal{R}, \quad i = 1, \dots,$$

with the interpretation that c_i is period i consumption given the payoffs realized up to that period. The agent chooses a consumption plan from a non-empty subset \mathcal{C} which we interpret as the set of feasible plans. The specification of \mathcal{C} will vary with the problem considered.

Given a random payoff stream $f_i, i = 1, \dots$, an agent with belief μ solves:

$$\max_{c \in \mathcal{C}} E_\mu E_\theta \sum_{i=1}^{\infty} \delta^i u(c_i). \quad (2)$$

Let $c^*(\mu, f, \mathcal{C})$ be a solution to this problem. The *indirect utility* function is the value $V(\mu, f, \mathcal{C})$ of the above problem. That is,

$$V(\mu, f, \mathcal{C}) = E_\mu E_\theta \sum_{i=1}^{\infty} \delta^i u(c_i^*(f_1, \dots, f_i)).$$

⁷ For example Hansen and Sargent (2001) write: “Knight (1921) distinguished risky events, which could be described by a probability distribution, from a worse type of ignorance that he called uncertainty and that could not be described by a probability distribution. [...] A person behaving according to Savage’s axioms has a well-defined personal probability distribution. [...] Savage’s system undermined Knight by removing the agents possible model misspecification as a concern of the model builder.”

We suppress references to f and \mathcal{C} when they are clear from the context. In many applications, \mathcal{C} has a recursive structure, and the impact of past payoff realizations can be summarized by a vector of state variables.

A basic intuition is that agents prefer prospects with known probabilities to ones where the probabilities are unknown. In fact, Knight's original motivation was that economic profits are paid to bearing uncertainty. The first criterion to measure the impact of uncertainty is in terms of the expected utility improvement if the agent knew the true parameter:

$$E_{\mu} V(\theta) - V(\mu). \quad (3)$$

This term is strictly positive if V is strictly convex in θ . It captures the value of information in helping the agent make better consumption decisions. We examine consumption smoothing under uncertainty in greater detail in Section 3. For now, we simply note that uncertainty has no impact when considering preferences over consumption streams, since there is no decision to be made. Formally, a constraint set \mathcal{C} *precludes intertemporal smoothing* if $c_i = f(s_i)$ for all i and s .

Proposition 2.1 *If \mathcal{C} precludes intertemporal smoothing then for every f , $V(\mu) = V(\bar{\theta}^{\mu})$.*

Precluding intertemporal smoothing means that the agent consumes his endowment, so indirect utility reduces to utility over consumption streams. The agent cannot, for example, open a checking account, store consumption goods, or put money under the proverbial mattress. In realistic economic environments, some degree of intertemporal smoothing is possible, and uncertainty potentially has an impact on behavior.

The second criterion to measure the impact of uncertainty is to use as a reference point the agent's utility under the average parameter:

$$V(\bar{\theta}^{\mu}) - V(\mu) \quad (4)$$

and define the *uncertainty premium* as:

$$u^{-1}(V(\bar{\theta}^{\mu}, f)) - u^{-1}(V(\mu, f)).$$

To illustrate these concepts, we focus on a concrete class of examples:

Example 1 (Deferred Consumption) *There are $n+1$ periods and $\delta = 1$. The set \mathcal{C} consists of a single consumption plan: $c_i = 0, i \neq n+1$ and $c_{n+1} = \sum_{i=1}^n f_i$.*

One interpretation of the utility of consumption in period $n+1$ is that it represents the indirect utility of consumption in subsequent periods with initial wealth given by the lump-sum payment $\sum_{i=1}^n f_i$. Many important problems fit this description, including assets that generate payoffs each period but pay the cumulative dividend at time $n+1$.⁸

The separation of consumption (period $n+1$) and payoff accumulation (periods $i = 1, \dots, n$) makes the problem tractable. Assuming $u(0) = 0$, indirect utility is simply:

$$V(\mu) = E_\mu E_\theta u\left(\sum_{i=1}^n f_i\right).$$

Note that the term (3) is zero, so that criterion cannot separate risk and uncertainty. The separation is possible under criterion (4). However, when u is strictly concave and $n > 1$, the distribution of the sum $\sum_{i=1}^n f_i$ has more variability under the uncertain belief μ than under the i.i.d. parameter $\bar{\theta}^\mu$. This suggests that an agent will prefer a project with known probability $\bar{\theta}^\mu$ to a project with uncertain probability represented by μ .

This intuition is confirmed by Halevy and Feltkamp (2005) in the case of $n = 2$. They consider an agent whose utility depends on the sum $f(s_1) + f(s_2)$ of two draws from an urn with fixed but unknown composition. They show that this creates sensitivity to uncertainty: a risk averse agent would prefer to bet on an urn with known composition $\bar{\theta}^\mu$ to an uncertain urn with subjective distribution μ , so (4) holds strictly. Halevy and Feltkamp (2005)'s important insight is that many real-world situations involve multiple draws and that utility may depend on the sum of these draws. They suggest that agents may develop heuristics that make them appear sensitive to uncertainty even in one-draw experimental settings.

⁸ For example, a start-up company which accumulates gains and losses to its value over a period of time, but investors are paid when the company is sold or goes public. Another example is a retirement portfolio that generates dividends each period, but the agent cares only about its value at the time of retirement.

Halevy and Feltkamp (2005)'s study shows that analyzing even the $n = 2$ case can be quite involved. We gain additional intuition by considering large n . Imagine a partnership where each partner owns (conveniently) a $\frac{1}{n}$ share of the final value of the firm. This normalization eliminates the effect of possible changes in risk attitude as n grows. The agent has a choice between two projects: one with known odds $\bar{\theta}^\mu$ and another with uncertain odds μ .

Proposition 2.2 *Suppose that u is strictly concave and μ is non-degenerate. Then for all sufficiently large n , the agent strictly prefers the project with known odds.*

This says that $V(\bar{\theta}^\mu) > V(\mu)$ and the uncertainty premium is strictly positive. With large n , the intuition is simple: the law of large numbers implies that for any θ there is high probability that the average is close to the expectation $E_\theta f$, hence the approximate equality:

$$E_\theta u\left(\frac{1}{n} \sum_{i=1}^{\infty} f_i\right) \simeq u\left(\frac{1}{n} E_\theta \sum_{i=1}^{\infty} f_i\right).$$

The distinction between parameters and subjective beliefs about parameters is key: while the variability implied by θ tends to average out, uncertainty is unaffected by n . We finally note that uncertainty is irrelevant in a one-period problem. With $n = 1$, $P^{\bar{\theta}^\mu}$ and P^μ induce identical distributions on outcomes, and are therefore indistinguishable. Our model is therefore inconsistent with behavioral anomalies that arise in one-period choice problems.

3 Uncertainty and Precautionary Behavior

In this section we turn to a richer model where consumption decisions and income realizations occur in each period. This standard consumption-saving setting makes it possible to examine the theoretical and empirical implications of uncertainty in a familiar context.

3.1 Model

In each period $t = 1, \dots$, the agent starts with a level of wealth w_t and belief μ_t which represents the *physical state* and *belief state* of the system, respectively. His consumption decision is a function of the state, $c_t(w_t, \mu_t)$.

The physical state evolves according to:

$$w_{t+1} = (1 + r)(w_t + f_t - c_t), \quad (5)$$

where c_t is consumption in period t and $r > 0$ is the net return on savings. Beliefs evolve according to Bayes rule. Note that the transition of the physical state does not depend on beliefs, and conversely. The transitions on wealth and beliefs define a feasible set \mathcal{C} . The value function (indirect utility) and the policy function are denoted $V(\mu, w)$ and $c^*(\mu, w)$.

To obtain an analytical expression for the evolution of consumption and wealth under uncertainty, we assume that the agent has an exponential (CARA) utility function $u(x) = -e^{-ax}$ and that $\delta = \frac{1}{1+r}$. Our analysis for the pure risk case is based on Caballero (1990), while the results for uncertainty appear new. See Carroll and Kimball (2008) for a recent survey.

The optimal solution to this problem must satisfy the Euler equation:

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

For the exponential utility, marginal utility is a constant multiple of the utility level. Using this fact, and iterating expectations, we have:

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

for every k . This implies (by taking u^{-1} of both sides) that consumption today equals the certainty equivalent of consumption at any future date. This in turn implies that the value function satisfies:⁹

$$V(\mu_t, w_t) = u(c_t^*).$$

3.2 Pure Risk

The evolution of consumption under risk and exponential utility is well-known in the literature on precautionary saving. We summarize these results for the benefit of the reader:

⁹ This follows from: $V(\mu_t, w_t) = (1 - \delta)E_{\mu_t} \left[\sum_{l=t}^{\infty} \delta^{l-t} u(c_l^*) \right] = (1 - \delta) \sum_{l=t}^{\infty} \delta^l u(c_t^*)$.

Proposition 3.1 *Given any parameter θ , the optimal consumption rule is:*

$$c^*(\theta, w) = \frac{r}{1+r} w + r^{-1} CE(rf|\theta) \quad (6)$$

where

$$CE(rf|\theta) \equiv u^{-1} \left(\int u(rf) d\theta \right)$$

is the certainty equivalent of the random permanent impact rf on consumption from the dividend implied by θ .

Furthermore, the evolution of consumption can be written

$$c_{t+1} = c_t + \Gamma(\theta) + r(f_t - E[f|\theta]) \quad (7)$$

where $\Gamma(\theta)$ is the risk premium of the variable rf , defined by

$$\Gamma(\theta) \equiv E(rf|\theta) - CE(rf|\theta) \geq 0. \quad (8)$$

Examining formula (7), we see that Γ is the drift of the consumption process (since the last term has mean zero). Once again, Γ equals the standard risk premium of the prospect rf . In particular, Γ is increased if the distribution on f is replaced by a mean-preserving spread. It is strictly positive as long as θ is not a point mass, since a certainty equivalent (under concave utility) is always less than expected value. It should be interpreted as due to “precautionary savings” since it would be absent if future dividends were known. The term $r(f_t - E[f|\theta])$ is a mean-zero random shock to consumption due to the wealth effect of the difference between the realized and expected payoff.

We can also observe that, by inspection of the CARA utility form, the term $r^{-1}CE(rf|\theta)$ which appears in (6) is in fact the certainty equivalent of the prospect f for an agent who has CARA utility with risk aversion of ra rather than a . Notice that this must be larger than $CE(f|\theta)$ whenever $r < 1$; an agent who is less risk averse has larger certainty equivalents. An equivalent statement is that for $r < 1$:

$$CE(rf|\theta) \geq rCE(f|\theta)$$

which is just the statement that there is less risk aversion at a smaller scale. The inequality is strict whenever f takes on more than one value.

3.3 Uncertainty

Here we let μ be a non-trivial distribution over the parameter θ . The following general form, stating that adding wealth simply results in additional interest income being consumed each period, follows intuitively from the fact that risk attitudes under CARA are invariant to the wealth level:

Lemma 3.2

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

In light of the lemma, it makes sense to use the notations

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

The next result shows that the value and consumption functions are convex in μ ; this can be interpreted to mean that given any uncertainty, resolution of this uncertainty will on average increase consumption. The last part states further that, for a fixed marginal distribution on tomorrow's return, we consume less under uncertainty than under pure risk. That is, there is additional precautionary savings due to the uncertainty about θ . To take an extreme example, we consume more if we know our returns will be i.i.d. 50-50, 0 or 1, forever, than if we think there is a 50% chance they will be 0 forever and a 50% chance they will be 1 forever.

Proposition 3.3 *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

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

The proof of the last statement makes it clear that there are two forces at work in the preference for the risk represented by $\bar{\theta}^\mu$ over the uncertainty represented by μ . First, there is the value of information; it is better to have uncertainty resolved at time 0, because it allows for superior consumption planning. Formally, this value of information is equivalent to convexity with respect to μ . Second, there is the hedging motive of having many independent gambles, which is valuable when consumption can be smoothed. ¹⁰

Lemma 3.2 implies that

$$\begin{aligned} c_{t+1} &= c_t + \frac{r}{1+r} (w_{t+1} - w_t) + c^*(\mu_{t+1}) - c^*(\mu_t) \\ &= c_t + r(f_t - c^*(\mu_t)) + c^*(\mu_{t+1}) - c^*(\mu_t) \end{aligned}$$

To interpret the evolution of consumption under uncertainty, it is useful to write it in a way that parallels that under risk, (7). Let

$$\Gamma(\mu_t) \equiv r(E[f_t] - c^*(\mu_t)),$$

and note that this reduces to (8) when μ has support θ . This term is certain as of time t and represents an upward drift in consumption similar to what we have seen under risk.

Rearranging terms, the evolution of consumption can be written as:

$$\begin{aligned} c_{t+1} &= c_t + \Gamma(\mu_t) + r [f_t - E_{\mu_t}[f_t]] \\ &\quad + [E_{\mu_t}[c^*(\mu_{t+1})] - c^*(\mu_t)] + [c^*(\mu_{t+1}) - E_{\mu_t}[c^*(\mu_{t+1})]]. \end{aligned} \quad (9)$$

The first line in (9) closely resembles behavior under risk. The second line includes a new source of upwards drift and a new mean-zero random shock to consumption. Both terms are due to the resolution of uncertainty and therefore absent under risk. The upward drift in consumption is due to the resolution of uncertainty which implies less precautionary savings in the future. This term is positive because Bayesian updating implies that $\mu_t = E_{\mu_t}[\mu_{t+1}]$, and by Jensen's inequality and the previous proposition, this means that $E_{\mu_t}[c^*(\mu_{t+1})] \geq c^*(\mu_t)$, with strict inequality whenever μ_t is not a point mass. The new source of randomness (second term in the second line) is due

¹⁰ Formally, this is concavity with respect to θ as shown in Proposition A.1.

to the persistent effect of the random resolution of uncertainty – that is, the current dividend is informative about future dividends. This was not present under risk.

4 Empirical Implications of Uncertainty

How is uncertainty different from other situations of imperfect information? In a Bayesian setting, both “measurable risk” and an “unmeasurable uncertainty” correspond to the agents’ probabilistic belief about his environment. We will argue, however, that risk and uncertainty have sharply different empirical content.

4.1 Uncertain Beliefs are Untestable

A natural intuition is that the probabilities of uncertain events are, in a sense, subjective. They are, in Keynes’ words, events “about [which] there is no scientific basis on which to form any calculable probability whatever.” This has important implications for disagreement and belief heterogeneity. One would expect most people to agree on the probability of heads in a fair coin toss, but not on the probability of the outcome of a presidential election or the unraveling of a monetary union. We formalize the distinction between objective risk and subjective uncertainty in terms of statistical tests with suitable properties. We show that objective statistical tests cannot determine whether uncertain beliefs are right or wrong.

Consider a simple setting where only two parameters, $\theta_1 \neq \theta_2$, are relevant. Let Δ be the set of beliefs μ with $\mu(\theta_1) + \mu(\theta_2) = 1$ and $\mu(\theta_i) > 0, i = 1, 2$. Interpret a function

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

as a statistical test, where $T(\theta_m, s^\infty) = 1$ indicates that the infinite sequence of observations s^∞ confirms the hypothesis that the true data generating process is θ_m , while $T = 0$ means that s^∞ is inconsistent with θ_m . As an idealization of what is in principle testable, consider asymptotic tests that

make use of the entire infinite sequence of data.¹¹

With unlimited data, it is natural to require a test to have the following properties:

1. T is free of Type I error on parameters:

$$P^{\theta_m} \{T(\theta_m, s^\infty) = 1\} = 1 \quad \forall m.$$

2. T identifies the true parameter:

$$P^{\theta_m} \{T(\theta_{m'}, s^\infty) = 1\} = 0 \quad \forall m \neq m'.$$

A simple test with these properties is one that compares θ_m with the empirical frequencies along the sequence of data. In the absence of uncertainty, such a test can be used to objectively determine whether the belief that the true parameter is θ_m is right or wrong.

This objective measurement of the correctness of beliefs is not possible under uncertainty. To make this precise, fix a test T with the above properties and let $\hat{T} : \Delta \times S^\infty \rightarrow \{0, 1\}$ be any Type I error-free extension of T (that is, $P^\mu \{T(\mu, s^\infty) = 1\} = 1$ for all $\mu \in \Delta$). Requiring small Type I error (or zero, in the asymptotic limit) reflects the priority given to controlling this type of error in statistical practice.

Proposition 4.1 *Let μ, μ' be any pair of beliefs in Δ and \hat{T} any test with the above properties. Then*

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

Consider agent i who believes the economy evolves according to μ_i . This agent is therefore convinced that any other agent j who disagrees with him, $\mu_j \neq \mu_i$, must be wrong. The proposition says that this agent must also believe that there is no Type I error-free test that can objectively determine which belief is correct.

Proof: We have

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

¹¹ Footnote 12 below discusses finite tests.

This implies that $P^{\theta_1}\{T(\mu', s^\infty) = 1\} = P^{\theta_2}\{T(\mu', s^\infty) = 1\} = 1$. From this it follows that:

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

The other equality is proved similarly. ■

Parameters are objective in the sense that one can devise a powerful statistical test that verifies the value of the true parameter. Uncertainty about parameters, by contrast, is subjective because any unprejudiced test of these beliefs has no power, and thus cannot be used to refute or confirm these beliefs. Uncertainty is different from risk not because agents do not use probabilities in making decisions, but because it is difficult to devise objective criteria to test them.¹²

4.2 Inference under Uncertainty

The discussion above suggests that uncertainty is a state of beliefs whose correctness cannot be objectively tested. On the other hand, uncertain beliefs influence decisions and can have important observable implications. For concreteness, we use a simple parametrized version of the consumption-saving model to show that incorporating uncertainty in empirical studies raises special challenges.

4.2.1 The Rational Expectations Benchmark

Empirical methods in dynamic economic models usually assume that agents have rational expectations, *i.e.*, beliefs coincide with the observed empirical frequencies. In the stylized saving-consumption model of Section 3, we define *the rational expectations benchmark* as an agent's belief that coincides with the observed empirical frequencies of the exogenous income process f .

¹² A test \hat{T}_n based on n observations cannot distinguish between knowing θ_m with certainty and a belief μ that puts arbitrarily large weight on θ_m . The argument in the proposition can be modified to cover finite tests. We show this with an example: Suppose that $\mu(\theta_1) = \mu(\theta_2) = 0.50$ and $P^\mu\{\hat{T}_n(\mu, s^\infty) = 1\} \geq 1 - \epsilon$, where ϵ is a small positive number representing Type 1 error. This says that \hat{T}_n has small Type I error at μ . Using the same argument as the proposition, it is easy to show that $P^\mu\{\hat{T}_n(\mu', s^\infty) = 1\} \geq 1 - 2\epsilon$.

Take the perspective of an outside modeler (an econometrician) who combines restrictions inspired by an economic model with past observations of income realizations to estimate parameters such as θ and a . If the sequence of observations is long enough, the modeler’s estimate, denoted $\hat{\theta}$, will be close to the true parameter $\bar{\theta}$ with high probability. Roughly, rational expectations econometrics estimates the remaining model parameters by assuming that the agent’s decisions are optimal with respect to the belief $\hat{\theta}$. These “cross-equation restrictions” are a powerful empirical tool that eliminates expectations as a free variable.

For our numerical example, assume that income takes just two values, $f \in \{0, 1\}$, the interest rate is $r = 0.01$, and that the true parameter is $\bar{\theta} = 0.50$. After a long sequence of observations, the modeler’s point estimate $\hat{\theta}$ will be close to 0.50 with high probability. Since the inference of the econometrician is not our main concern here, assume for convenience that his estimate is not subject to sampling error: $\hat{\theta} = \bar{\theta}$. The simplest way to impose rational expectations is to assume that the agent’s belief μ puts unit mass on the sample estimate:

$$\mu(\hat{\theta}) = 1.$$

Rationality implies that past choices are optimal given this belief.

In this example, the hypothesis of rational expectations is a key identifying assumption that pins down beliefs to observations. This leads the modeler to assume that consumption evolves according to the autoregressive equation (7):

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

where $r(f_t - E[f|\hat{\theta}])$ is a mean-zero i.i.d. disturbance and $\Gamma(\hat{\theta})$ is a drift term.

This model has a number of concrete empirical implications that can be tested against income and consumption data. For example, the model predicts that the drift term is $\Gamma(\hat{\theta})$, where $\hat{\theta}$ is estimated directly from observed income. Letting $\hat{\alpha}$ denote the coefficient in a simple linear regression of $c_{t+1} - c_t$ on a constant term, the model’s null hypothesis is that $\hat{\alpha} = \Gamma(\hat{\theta})$, and that the residuals are generated by the i.i.d. process $r(f_t - E[f|\theta])$. In particular, the residuals must be serially uncorrelated and homoskedastic.

4.2.2 Consumption under Uncertainty

The assumption that the agent’s beliefs equal the empirical distribution implicitly assumes that: (1) the agent observed a long sequence of past realizations of the process; *and* (2) the agent believes that the future will be similar to the past, in the sense that the same parameter that governed past income will continue to govern future income. Uncertainty, in the form of a non-degenerate μ , may represents a “belief shock” where the agent questions the stability of the parameter governing his environment and the relevance of past data to estimating the new parameter.

Under uncertainty, consumption evolves according to (9):

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

The belief state μ_t changes with the arrival of new information plays a crucial role in the decision problem. Under risk, by contrast, beliefs are not updated since the agent is assumed to know the true parameter $\bar{\theta}$.

Qualitatively, the introduction of uncertainty results in an additional drift term and a new mean-zero disturbance. However, the evolution of consumption itself cannot be derived analytically. To better understand the impact of uncertainty, we simulate the model numerically. Continue to assume that the true parameter is $\bar{\theta} = 0.50$ and that it is estimated without error by the econometrician after observing a long sequence of income realizations. The agent, who lacks this knowledge *ex ante*, quantifies his uncertainty with a uniform prior μ over $[0,1]$. Note that the agent is convinced that the parameter is ‘on average’ equal to $\bar{\theta}^\mu = 0.50$, in the sense that he believes the marginal on income in any period to be 50/50.

To give a sense of how consumption evolves under risk ($\bar{\theta}^\mu = 0.50$) and uncertainty (μ), four samples are shown in Figure 1. In each sample, an income realization is drawn under the true parameter $\bar{\theta} = 0.50$. Given income, we compute the optimal consumption under risk and uncertainty, using (7) and (9) respectively. In each of the four samples, consumption under uncertainty is lower in period 1, has more upward drift, and is more volatile.

Figure 2 summarizes the behavior of consumption averaged over 10,000

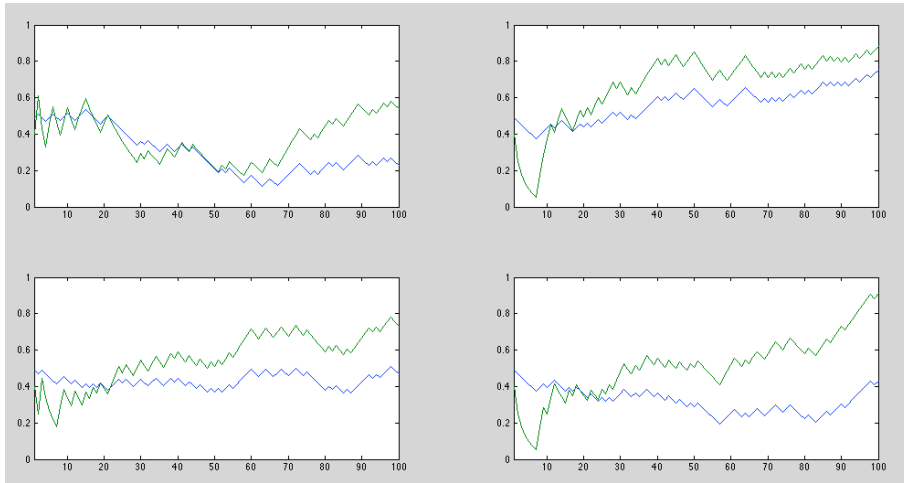


Figure 1 – Four Sample Paths of Consumption under Risk and Uncertainty

In each subplot, both lines come from the same sequence of draws, with $\theta = 0.50$. The blue line represents the consumption path under risk (known θ), while the green line represents consumption under uncertainty (the agent's has uniform prior over θ).

samples. Under risk (assuming known $\bar{\theta} = 0.50$), we have $c_{t+1} - c_t = \Gamma(\theta)$, so c_t increases linearly from an initial value c_1 that depends on θ, a , and r . An outside observer who uses past income realizations to obtain an estimate $\hat{\theta} = 0.50$ will expect consumption, averaged over samples, to evolve according to this linear relationship.

Under uncertainty, consumption starts at a lower level. This reflects the increased precautionary savings motive due to the agent's uncertainty about the true value of θ .¹³ Under the assumption that the data is generated by the parameter value $\bar{\theta} = 0.50$, consumption rises, eventually exceeding RE consumption. The reason is that the agent, who accumulated greater wealth initially, has more dividend income.

Take again the perspective of a modeler whose model incorporates the hypothesis that the agent's belief is equal to the sample estimate, $\hat{\theta}$. If

¹³ Precautionary saving is also present under risk. The difference in period 1 consumption in the figure is the *additional* precautionary saving due to uncertainty.

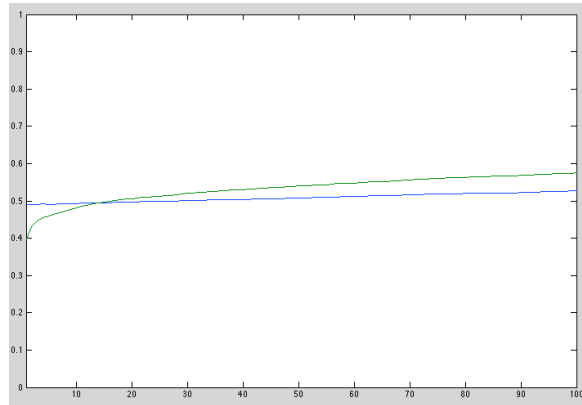


Figure 2 – Average Consumption under Risk and Uncertainty

*The straight line is consumption under risk averaged over 10,000 samples.
The concave line is average consumption under uncertainty.*

the agent is uncertain about θ , his behavior will evolve according to (9) instead, and the model is misspecified. The modeler will find, among other things, that residuals are serially correlated and heteroskedastic. As evident from Figure 2, serial correlation is a consequence of the steeper consumption growth resulting from a more aggressive precautionary saving early on.

4.3 Discussion and Related Literature

Lewellen and Shanken (2002) and Weitzman (2007) introduce learning-based models to explain asset pricing anomalies in terms of parameter-uncertainty. A major focus of these papers is the equity premium puzzle, namely that, for plausible levels of risk aversion, asset prices seem to imply a volatility in payoffs that is significantly higher than their empirical volatility (estimated, say, from the standard deviation of past returns).¹⁴ These authors argue that a possible source of this and other related puzzles is the econometric identification assumption that investors' beliefs about the future volatility of income and consumption are equal to these variables' historical volatility.

¹⁴ Equivalently, investors' risk aversion implied by observed asset prices is an order of magnitude higher than what is considered reasonable.

As noted earlier, this rational expectations assumption implicitly requires agents to believe that the parameter generating future outcomes is the same as the one that generated past data.

Although we do not consider asset pricing in this paper, similar effects appear in the precautionary saving model. Under our i.i.d. assumption, payoffs are generated by one true parameter $\bar{\theta}$, yet an agent who does not know its value will subjectively expect greater volatility in future income. This uncertainty may be the result of the agent fearing that a structural change in the income process might have occurred, shaking his confidence that the future will be similar to the past. This uncertainty can have important consequences for the behavior of endogenous variables but, as shown earlier, it may be difficult to objectively demonstrate through statistical tests that such uncertain beliefs are paranoid or irrational.

In Weitzman (2007)'s model, the data is generated by i.i.d. draws from a normal distribution with known mean but unknown standard deviation. An agent's subjective uncertainty about the standard deviation implies that his belief is a Student t distribution which, unlike the normal component distributions, has a fat tail. Weitzman shows that agents' concern about this tail risk can lead to significant implications for asset pricing. There is a parallel with what we do: in our model, a subjective distribution over i.i.d. parameters is not an i.i.d. distribution, while in Weitzman's model uncertainty over standard deviations of normals implies a subjective belief that is not itself normal. A second difference is that Weitzman does not consider intertemporal smoothing.

There is an extensive literature on precautionary savings. That literature mainly focuses on what we call risk (agents' beliefs coincide with the ergodic probabilities). An important exception is Cogley and Sargent (2008) who consider a model of precautionary savings with parameter uncertainty. Cogley and Sargent (2008)'s main focus is on whether behavior based on anticipated utility approximates well the full Bayesian solution in simulations under power utility. They assume an income process generated by a Markov matrix with unknown coefficients. The agent's prior on these coefficients is a Dirichlet distribution, a tractable modeling assumption that makes it easy to represent the extent to which an agent trusts his prior. The

uniform distribution used in our simulation reported earlier is a binomial Dirichlet (*i.e.*, Beta) distribution with parameters $n_0 = n_1 = 1$. As with Cogley and Sargent (2008), we can model the agent's uncertainty by varying the parameter of his prior. Pure risk is then the limiting case where n_0 and n_1 approach infinity. We may interpret this, as they do, that the agent has observed a long history of the process and that he believes that the future will look like the past. The advantage of CARA utility is that there are no wealth effects, making the model more tractable. This allowed us to derive equation (9) which showed clearly the relation between behavior under risk and uncertainty in an explicit equation and proved that uncertainty causes increased precautionary savings and increased volatility. While power utility is considered more realistic, the differences are very minor for any modest changes in wealth, and so we consider the ability to write down the solution more explicitly worthwhile. Of course, we still had to solve the model numerically to calculate $V(\mu)$, but the theory behind the calculation is made much simpler by the separability of wealth in the CARA model.

A Appendix

Proof of Proposition 2.2:

$$V(\mu) = E_\mu E_\theta u\left(\frac{1}{n} \sum_{i=1}^n f_i\right) \quad (10)$$

$$\simeq E_\mu u\left(\frac{1}{n} E_\theta \sum_{i=1}^n f_i\right) \quad (11)$$

$$\simeq E_\mu u\left(\frac{1}{n} E_\theta f\right)$$

$$\simeq E_\mu V\left(\theta, \frac{1}{n} f\right)$$

$$< V(\bar{\theta}^\mu),$$

where the last inequality follows from the strict concavity of u .

Proof of Proposition 3.1: We first note that since we are optimizing a concave function, it suffices to show that the Euler and transversality conditions are satisfied. Applying u^{-1} to both sides of $u(c_t^*) = E_t[u(c_{t+k}^*)]$ and substituting in (5), we can express the evolution of wealth as:

$$\begin{aligned} w_{t+1} &= (1+r) \left[w_t + f_t - \frac{r}{1+r} w - r^{-1} CE(rf|\theta) \right] \\ &= w_t + (1+r)f_t - (1+r)r^{-1} CE(rf|\theta), \end{aligned}$$

implying

$$\begin{aligned} c_{t+1} &= \frac{r}{1+r} [w_t + (1+r)f_t - (1+r)r^{-1} CE(rf|\theta)] + r^{-1} CE(rf|\theta) \\ &= \frac{r}{1+r} w_t + r f_t - CE(rf|\theta) + r^{-1} CE(rf|\theta) \\ &= c_t + r f_t - CE(rf|\theta). \end{aligned}$$

Under CARA, certainty equivalents are independent of wealth, so the last equation implies that c_{t+1} has a time- t certainty equivalent of c_t , which we argued earlier is equivalent to the Euler equation.

To get the formula for the evolution of consumption, note that θ does not change in time (this is the key property of lack of uncertainty) so we can deduce from equation (6) that

$$c_{t+1} - c_t = \frac{r}{1+r} (w_{t+1} - w_t)$$

after which substituting for w_{t+1} according to equation (5) and c_t from equation (6) gives (7). It shows that consumption, which is linear in wealth, can grow only linearly so long as f is bounded, so wealth also grows at most linearly also and the transversality condition $\lim_{t \rightarrow \infty} (1+r)^{-t} w_t$ is satisfied with probability 1. ■

Next, we establish the concavity of the value function under risk. Note that for any lottery over two outcomes, we have:

$$u^{-1}(\lambda u(r f_1) + (1 - \lambda)u(r f_2)) \geq r u^{-1}(\lambda u(f_1) + (1 - \lambda)u(f_2)) \quad (12)$$

This observation, along with the previous proposition, allows us to show:

Proposition A.1 *For known θ ,*

1. *Optimal consumption is convex in θ .*
2. *Value is concave in θ for any $r < 1$.*

To illustrate the meaning of this proposition, assume that under distribution θ_0 an asset returns 0 for sure, under θ_1 it returns 1 for sure, and thus under $\theta_2 = .5\theta_0 + .5\theta_1$ the asset returns 0 or 1 each period according to a series of i.i.d. coin flips. With zero wealth, optimal consumption under θ_0 is simply 0, with value $u(0)$, and under θ_1 is 1, with value $u(1)$. The first point then states that consumption under θ_2 is less than .5; this is because there is a small precautionary saving to mitigate against the risk of the next coin flip. However, value under θ_2 is greater than the average of $u(0)$ and $u(1)$; this is driven by consumption smoothing, as θ_2 ensures many independent coin flips and entails much less long-run risk than a lottery between θ_0 and θ_1 . There are two countervailing forces on the value, as taking a convex combination of θ s increases the need for precautionary savings but mitigates against long-run risk. The CARA utility allows us to say that the second force dominates for $r < 1$. When $r \ll 1$ as usual, the second force is much stronger and it is reasonable to think that the CARA form is not essential to this claim. It is interesting to note that if we allow $r > 1$, the precautionary-savings effect dominates, as discounting is fast and long-run risk is unimportant (mostly the current period matters.)

Proof of Proposition A.1: Applying Proposition 3.1, the first item reduces to the statement that the certainty equivalent of a compound lottery is less than the average of the certainty equivalents of the nodes. That is,

$$\begin{aligned}
CE(rf|\lambda\theta_1 + (1 - \lambda)\theta_2) &= u^{-1}E[u(rf)|\lambda\theta_1 + (1 - \lambda)\theta_2] \\
&= u^{-1}(\lambda E[u(rf)|\theta_1] + (1 - \lambda)E[u(rf)|\theta_2]) \\
&= u^{-1}(\lambda u(CE(rf|\theta_1)) + (1 - \lambda)u(CE(rf|\theta_2))) \\
&\leq \lambda CE(rf|\theta_1) + (1 - \lambda)CE(rf|\theta_2)
\end{aligned}$$

where the last step follows from u being increasing and concave (take u of both sides to get the definition of concavity). Along with Proposition 3.1 this yields the first statement.

For the second statement, use $V(\theta)$ as shorthand for $V(\theta, 0)$. By earlier results,

$$V(\theta) = u(r^{-1}CE(rf|\theta)) = u(r^{-1}u^{-1}E(u(rf)|\theta)).$$

Now observe that

$$\begin{aligned}
V(\lambda\theta_1 + (1 - \lambda)\theta_2) &= u(r^{-1}CE(rf|\lambda\theta_1 + (1 - \lambda)\theta_2)) \\
&= u(r^{-1}u^{-1}(E[u(rf)|\lambda\theta_1 + (1 - \lambda)\theta_2])) \\
&= u(r^{-1}u^{-1}(\lambda E[u(rf)|\theta_1] + (1 - \lambda)E[u(rf)|\theta_2])) \\
&= u(r^{-1}u^{-1}(\lambda u(ru^{-1}V(\theta_1)) + (1 - \lambda)u(ru^{-1}V(\theta_2)))) \\
&\geq \lambda V(\theta_1) + (1 - \lambda)V(\theta_2)
\end{aligned}$$

where the last step uses (12). This shows concavity of V in θ for $w = 0$. For general w , equations (3.1) and (6) show that

$$V(\theta, w) = e^{-a\frac{r}{1+r}w} V(\theta)$$

which implies the result. ■

Proof of Lemma 3.2: The Bellman equation tells us

$$V(\mu, w) = (1 - \delta)u(c^*(\mu, w)) + \delta E[V(\mu|f, (1 + r)(w - c^*(\mu, w) + f))]$$

and we showed earlier that $V(\mu, w) = u(c^*(\mu, w))$, so the Bellman equation reduces to

$$u(c^*(\mu, w)) = E[u(c^*(\mu|f, (1+r)(w - c^*(\mu, w) + f)))]$$

and when we substitute the form in the lemma, this reduces algebraically to the Bellman equation for $w = 0$. This proves the result. ■

Proof of Proposition 3.3: The first inequality (in its weak form) is simply the statement that information about θ has a non-negative value. Since the utility being optimized is strictly concave, optimal plans are unique and it is strict whenever plans differ under the two beliefs.

The second inequality follows from the first, given our earlier result that $V(\mu, w) = u(c^*(\mu, w))$ and the fact that u is increasing and concave. (Assume it is false, apply u to both sides, then apply concavity to contradict the first inequality.)

The final statement follows from expressing μ as a convex combination of point masses θ_i , applying convexity of V with respect to μ , then applying concavity of V with respect to θ (Proposition A.1). ■

References

- BEWLEY, T. (1986): “Knightian Decision Theory: Part I,” Cowles Foundation Discussion Paper no. 807.
- CABALLERO, R. (1990): “Consumption Puzzles and Precautionary Savings,” *Journal of Monetary Economics*, 25(1), 113–136.
- CARROLL, C., AND M. KIMBALL (2008): “Precautionary Saving and Precautionary Wealth,” *The New Palgrave Dictionary of Economics*, 6, 579–584.
- COGLEY, T., AND T. SARGENT (2008): “Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making,” *International Economic Review*, 49(1), 185–221.
- GILBOA, I., AND D. SCHMEIDLER (1989): “Maxmin Expected Utility with Nonunique Prior,” *J. Math. Econom.*, 18(2), 141–153.
- HALEVY, Y., AND V. FELTKAMP (2005): “A Bayesian Approach to Uncertainty Aversion,” *Review of Economic Studies*, 72(2), 449–466.
- HANSEN, L., AND T. SARGENT (2001): “Acknowledging Misspecification in Macroeconomic Theory,” *Review of Economic Dynamics*, 4(3), 519–535.
- KEYNES, J. M. (1937): “The General Theory of Employment,” *The Quarterly Journal of Economics*, 51(2), 209–223.
- KNIGHT, F. (1921): *Risk, Uncertainty and Profit*. Harper, New York.
- LEROY, S., AND L. SINGELL (1987): “Knight on Risk and Uncertainty,” *The Journal of Political Economy*, 95(2), 394–406.
- LEWELLEN, J., AND J. SHANKEN (2002): “Learning, Asset-Pricing Tests, and Market Efficiency,” *The Journal of finance*, 57(3), 1113–1145.
- PÁSTOR, L., AND P. VERONESI (2009): “Learning in Financial Markets,” *Annu. Rev. Financ. Econ.*, 1(1), 361–381.
- WEITZMAN, M. (2007): “Subjective Expectations and Asset-Return Puzzles,” *The American Economic Review*, 97(4), 1102–1130.