

Preference Representations for Catastrophic Risk Analysis

Richard E. Ericson*, Jamie L. Kruse†

June 2013

Abstract

This note surveys some behavioral models of preference relations applicable for analysis of decisions in the face of catastrophic risk. Catastrophic risk is characterized, and the implications of various representations explored in simple, illustrative examples. An argument is presented for the more general applicability of “variational preferences” for the analysis of behavior and decision making in the face of catastrophic risk.

1 Introduction

Catastrophic events, such as major hurricanes and earthquakes, tornados, floods and wild fires, have, over the past decade, become that focus of a growing body of economic literature.¹ These high (negative) impact, (very) low probability, events (“tail events,” in Nordhaus, 2012) appear to be imposing growing costs on society, raising policy issues that call for economic analysis. Such analysis must rest on an understanding of the human behavioral response to such events and the decisions taken prior to, and in anticipation of, their possible occurrence, in order to avert them and/or mitigate their impact. The need for such analysis is given some urgency by the perception that global climate change has rendered such events increasingly likely and volatile, and thus they have the potential to impose unprecedented and unpredictable damage when they occur. Indeed, at the global level, climate change itself may evolve into a massive catastrophe, threatening the ability of the earth to support modern civilization and standards of living, or even life itself.²

1.1 The Problem

The economic analysis of such potential catastrophes faces several analytic challenges. Most fundamental is the true *uncertainty* surrounding both the likelihood, the potential magnitudes, and the timing of these rare events. While there is limited frequency data for some such events (e.g. 100-year floods, EF-5 tornadoes, Cat-5 Hurricanes), consequences of individual occurrences are often unprecedented (Hurricanes Katrina, Sandy), and others without precedent in recorded history (catastrophic asteroid hit, catastrophic climate change). Thus the distribution, the likelihood, of such events is truly unknown, and their occurrence may be arbitrarily far in the future, or ‘the day

*Department of Economics, East Carolina University

†Department of Economics, East Carolina University

¹This work is both reflected in, and stimulated by the 30 October 2006 Stern review on the economics of climate change. See references and discussion in, for example, Nordhaus (2007) and Weitzman (2007, 2009a). Also see Parson’s (2007) review of Posner’s influential book.

²Other massive catastrophes have also begun to receive some attention, e.g. the impact of a sufficiently large asteroid, capable of eliminating much, or most, life on earth. See Chichilnisky, Eisenberger (2010).

after tomorrow’. This fundamental uncertainty compounds the better understood risks associated with any of the potential scenarios that might unfold in the face of these events, including scenarios in their absence.

As such events may only ‘occur’ well into the future, economic analysis should consider the interaction between uncertainty and the temporal dimension. Uncertainties regarding capabilities, technologies and resources, constraints and preferences increase as the horizon lengthens. And the kinds of decisions to be made, the available choices, will change over time, both exogenously and driven by the anticipation of the catastrophic events of interest.

Finally, behind any economic analysis must stand a model of human behavior, a model of the objectives, preferences, and beliefs of those making decisions either in anticipation or under the impact of the catastrophic event, or that might influence its likelihood and/or impact. In most analysis this involves using an “expected utility” representation of suitably monotonic (non-satiated) preferences based on a clear understanding/knowledge of both the outcomes and consequences of decisions and their probabilities/likelihoods. This models decision making in the face of well understood risks. While assuming such knowledge is a reasonable approximation in much policy analysis, it is questionable for extremely rare, unprecedented, and/or distant future events, particularly those of a catastrophic nature. Further, the standard model uses exponential discounting of (expected) future returns in evaluating the consequences of decisions, which effectively eliminates consideration of any events/consequences sufficiently far removed in time.³

The problem of modeling decisions, and their driving preferences, in the face of catastrophic risks, integrates all the other, logically separate, challenges. For any model of decision making must incorporate not only tastes and preferences with regard to outcomes, their uncertainties, and their timing, but also knowledge about constraints, possible actions, and consequences, and beliefs about the relevant likelihoods of both major and minor events and changes in the constraints, technologies, threats and opportunities that may be faced. Here we explore only a part of that challenge, considering some analytic models of preference representation that we believe could be useful in analyzing human behavioral response to catastrophic events. We focus only on atemporal decision models, where ‘values’ are appropriately discounted, and the central question is: How do agents deal with fundamental, poorly defined uncertainties/risks?

1.2 Some Modeling Issues

There are a number of issues that should be kept in mind when developing an adequate model of decision making in the face of catastrophic risk. We highlight these issues not because we will or can deal with them adequately in this paper, but because they are important to interpreting our analysis, and will have to be dealt with in developing a fully adequate model of decisions in the face of such risks.

First, we consider ‘catastrophic’ risk to be more than the low-probability likelihood of extremely high losses. To be truly ‘catastrophic’ the extremely high losses/damage must be wide spread, impacting large numbers, and potentially causing a fundamental change in the socioeconomic system, in the ‘way the world works’. In such a (rare) situation, not only agent expectations, but fundamental understandings about technologies, constraints, and how actions map into consequences, as well as the preferences of the agent, are apt to be different from the ‘normal world’. And how all these will change is highly uncertain, indeed largely unknown and not considered, as agents make decisions *ex ante* with only the vaguest understanding of the circumstances with which they must deal in that rare/ unprecedented event. Thus catastrophic risk relates to a situation of fundamental

³for example, a 3% discount reduces \$1 million to \$493.10 in a quarter century. A 1% discount requires 757 years to achieve the same degree of devaluation, while a 5% discount requires only 14.85 years.

uncertainty, where the likelihood of outcomes and consequences is truly unknown, and the decision makers' beliefs and 'tastes' for this fundamental uncertainty become critical components of any decision model.

Another issue that must be faced is the purpose of the modeling exercise, the objective of the analysis. A model of decisions in the face of catastrophic risk might be either 'descriptive' or 'prescriptive', i.e. either capture actual human decisions and behavior, including behavioral deviations from 'economic rationality', or provide a rational framework for making 'best' decisions or policies (social rationality). This is closely related to the 'subject' of the analysis: Who's preferences are being modeled? If we are studying merely extreme risks, with some observed, if low, frequency (measurable probability) of occurrences, then expected utility with appropriate risk and ambiguity preferences, and appropriate discounting, provides a useful decision tool; it reflects a 'social rationality', even if it does not capture the behavior of (some/many) individuals. For public policy analysis and decisions, we require intertemporally consistent, equitable, and probabilistically sophisticated, expectationally 'objective' benefit-cost analysis. Such analysis should incorporate fundamental uncertainty about what is known and what will be known about the 'risk' and its consequences, as well as the fundamental uncertainty about future 'technology' and consequent 'rates of return', even if the appropriate pure discount rate is zero. We believe this is too much to assume for the analysis of individual behavior. Indeed, even for analysis of individual behavior in the face of moderate probability ('frequent' events), the model must consider individual time preference and 'behavioral' regularities, often violating standard axioms, such as: ignoring, or significantly overestimating ('fear factor'), very 'small' probabilities; underestimating, or accepting as 'certain', probabilities near one; hyperbolic or other non-exponential (time inconsistent) discounting; etc.

If a catastrophic event is of extremely small probability (rare event), the properties of the lower tail of distribution become significant for evaluation of 'measurable' events, as do (possibly non-measurable) 'tail events' for non-integral evaluation functionals. The standard analysis alluded to above is adequate for thin-tailed distributions (those dominated by Gaussian) over measurable events. However, uncertainty/ambiguity about even thin-tailed distributions can render the *decision-relevant distribution* 'heavy-tailed' (Weitzman, 2009a,b). Such distributions (power law; stable with $\alpha < 2$; regular variation at $-\infty$) lack finite moments of order above the first, and sometimes even the first moment, requiring direct evaluation of loss risk in tail, related to tail index of regular variation. And Chichilnisky (2000) has introduced catastrophes as 'non-measurable' events, requiring a non-integral evaluation functional, e.g. purely finitely additive functional, of 'events'.

Finally, any adequate model needs to be able to capture various agent (decision maker) responses to 'catastrophic risks'. This can range from ignoring to seriously exaggerating the dangers they pose, particularly when of very low, or unknown, (subjective) probability. 'Rational' behavior involves making choices ('acts') maximizing some value (minimizing some loss), given knowledge/beliefs/'fears' about likely consequences of those choices, including 'doing nothing' or 'waiting'. The valuations driving 'optimizing' choices are typically modeled in "preference representations" (utility functions) on a space of 'payoff relevant' outcomes, with 'beliefs' modeled by bounded, real-valued set functions (probabilities, capacities) on a space of payoff relevant 'events'. In this exercise, we believe that systematic deviations from modelled 'rational behavior' are best understood as a problem with the model — improper modelling of 'preferences' or 'beliefs' — requiring further development or new axiomatizations. If marginal utility is *unbounded* below ($x \downarrow$), the basic expected utility model suffers "tyranny of catastrophic risks" — outcomes with 'vanishingly small' probability, but an arbitrarily 'bad' result, dominate choice. The 'certainty equivalent' becomes the arbitrarily bad outcome implying an unbounded 'willingness to pay to avoid' (Buchholz, Schymura, 2010). On the other hand, if marginal utility is *bounded* below, the basic expected utility model eventually ignores risks, however catastrophic, of arbitrarily small probability ("black

swans”). This renders expected (subjective) utility analysis somewhat impotent in the face of truly catastrophic risks of the sort described above.

1.3 What We Do

This note is a first step toward our exploration of these issues. Here we focus on the impact of perceptions of likelihood and the decision maker’s ‘taste for uncertainty’. We present an extremely simple (2 or 3-outcome) example of a world with catastrophic risk, incorporating only the minimal elements required for such an example, and develop the implications of six different models of preferences in the face of such risk. This model world allows us to explore some of the consequences of the axioms differentiating the various decision models, as well as exploring those models’ implications for the behavior of agents when facing ‘catastrophic’ risks. We believe that this model world *may* provide basis for future ‘experimental’ work.

To capture agents’ beliefs and ‘taste’ for fundamental uncertainty, we will use models with ‘ambiguity’, models initially developed to resolve the behavioral “paradoxes” of the (subjective) expected utility model, including the Ellsberg and Allais/Dreze paradoxes, and the kind of behavioral anomalies addressed in prospect theory. These anomalies and paradoxes cast doubt on agents’ use of mathematically consistent subjective probabilities in making decisions, particularly when those probabilities are objectively unknown and/or near zero or one. We expect catastrophic risks to be particularly affected by these behavioral distortions, given their extreme rarity and frequently unprecedented nature when they do occur. Thus models in which the decision maker finds the likelihood/probabilities of events ‘ambiguous’ would seem particularly appropriate.

Below we present 6 such models, indicating their particular assumptions (axioms), and illustration their implications in our simple analytic example. The models we discuss are: (i) Choquet Expected Utility; (ii) Rank Dependent Expected Utility; (iii) Maximin Expected Utility; (iv) ‘Smooth’ Ambiguity Aversion (Second-Order Expected Utility); (v) Chichilnisky Model of “‘Sensitivity’ to ‘Rare’ Events”; and (vi) Variational Preferences. The final model we believe has the potential to become an ‘umbrella’ model, encompassing all of the others, depending on the specific situation modelled. The models share a common notation, summarized in the Appendix, where the full set of axioms used to derive the various representations is also presented.

2 Minimal Model of Catastrophic Risk

The idea of a ‘catastrophe’, rather than just an extremely bad outcome, must go beyond its rarity and unpredictability. It encompasses not only deep and wide spread (affecting a significant portion of the population) losses, but an aspect of fundamental change, an irreversibility of the resulting situation. Available resources and prior technologies can become sharply limited, standard production possibilities and means of interaction infeasible, requiring completely different ways of dealing with the new post-catastrophe world. And the resulting radical change in the ‘choice set’ suggests that preferences, including the taste for risk and uncertainty, are also apt to fundamentally change. While modeling in any detail such an unprecedented situation is perhaps an impossible task, it behoves us to incorporate whatever we can of the limited qualitative characteristics of such a world in our simple, reduced form representations. For any decisions in the face of the possibility, the risk, of such a catastrophic state must consider their consequences in that, as well as other, possible states.

With this in mind, we suggest that the following comprise a minimal set of essential components of any “Catastrophic Risk” model:

- A set of ‘states’ in which consumption/production opportunities, and hence welfare, are severely limited;
 - Very low probability (‘rare’) and very low utility (high loss);
 - Severe restriction of production possibilities, limiting ability to recover or raise consumption;
 - Uncertain timing and incidence;
- Decisions/‘Acts’, each with impact/consequences across all ‘states’;
 - Consequences of *ex-ante decisions/acts* are dramatically different in catastrophes than in non-catastrophic states;
- Shift in the ‘evaluation paradigm’ in the event of catastrophe: “state-dependent utility”⁴
 - *ex-ante* uncertainty/‘ambiguity’ with regard to likelihood of ‘catastrophic’ states;
 - *ex-post* different ordering over consequences, and attitude toward risk, in ‘normal’ and ‘catastrophic’ states.

In this note we will explore, for illustrative purposes, some of these aspects in an extremely simple 2- or 3-state example, aimed at capturing the implications of different approaches to modeling decision maker behavior in the face of uncertainty about the likelihood of the ‘catastrophe’. Thus it attempts to capture, if only in highly reduced form, the differences in consequences and utilities between ‘normal’ and ‘catastrophic’ states. The illustrative examples, and the decision models discussed, will use a common notation within the Anscombe-Aumann framework.

2.1 Illustrative Examples

Two examples are used to illustrate the implications of differing formulations/models of decision making in the face of catastrophic risk. They allow us to clearly show how differing assumptions about agent behavior impact the evaluation of a decision/action to be taken. The first is in a fully reduced form, where preferences/utility only reflect the risk attitude toward the occurrence of a catastrophic state. Hence it illustrates how differing assumptions about agents’ perceptions of the uncertainty affect their evaluations, the ‘certainty equivalent’ of the situation faced. The second allows (limited) exploration of how differing attitudes towards uncertainty, and differing risk attitudes in different states, might affect decisions.

2.1.1 2-outcome Example

In the simplest example, outcomes, $x(s)$, give the agent’s evaluation (expected utility) of the consequences of optimal decisions in state s , and her ‘utility’ function captures her attitude toward ‘state risk’. Let $x(s_0) = 0.5$; $x(s_1) = 12$ — the ‘catastrophic’ and ‘normal’ state outcomes, with probabilities (0.005, 0.995) respectively, and $u(x) = 1 + \frac{x^{1-\eta}}{1-\eta}$ reflects her constant relative risk aversion, with $\eta = 2$. Then the expected value of this situation is $E(x) = 11.943$, and the standard (expected utility) model gives $Eu(x) = 0.907$, with a certainty equivalent of 10.776. This provides a benchmark for comparison with other approaches to evaluating ‘state risk’ uncertainty. This example is used to illustrate differences in evaluations in the first five representations.

⁴While it is typically not possible to identify unique state probabilities and utilities in the derivable additively separable representation, we believe it useful to choose a specific representation that captures the intuitively plausible impact of the state on *ex-post* (Bernoulli) utilities.

2.1.2 3-outcome Example

This example is only used with respect to the sixth representation. It allows the introduction of decisions (Savage ‘acts’) in the face of uncertainty about the true state. Let the underlying states be $S = \{0, 1\}$, and let there be potentially different outcomes in each of the states, $X = \{0.5, 3, 9\}$, $X_0 = \{0.5, 3\}$, $X_1 = \{3, 9\}$. Here the probabilities of outcomes depend on the decisions/acts of the agents, and those will depend on the beliefs of the agent about the probability of the ‘bad’ state, $s = 0$. In contrast to the 2-outcome example, there is no given probability for outcomes, and hence expected value is not well defined. Let p be the probability of the underlying ‘bad’ state. Then, for example, letting $f := (f_0, f_1)$ be the distribution over outcomes in the two states induced by the act $f \in \mathcal{F}$, the ex-ante probability of $x = 3$ is $p \cdot f_0(3) + (1 - p) \cdot f_1(3)$.

To capture potential changes in preferences in the ‘catastrophe’ state (bad state, with worst outcome), we consider three Bernoulli utility functions, reflecting two different degrees of ‘risk aversion’:

$$u_0(x) = 1.15416 + \frac{x^{1-\gamma}}{1-\gamma} \text{ or } u_0(x) := v(x) = 0.5 + \frac{x^{1-\gamma}}{1-\gamma}, \quad \gamma = 3; \quad u_1(x) = \ln(x).$$

Note that the first u_0 and u_1 yield the same ex-post utility in the outcome ($x = 3$) that is common to both the catastrophe ($s = 0$), where it is ‘best’, and the normal state ($s = 1$), where it is worst. However, $v(3) < u_1(3)$, allowing illustration of the impact of distinctly lower welfare in the catastrophic state. Let $w_s(x)$ be the utility function in state s . Then the (subjective) expected utility (SEU) model yields evaluation of act f , $Ew(f) = \sum_x p \cdot f_0(x)w_0(x) + (1 - p) \cdot f_1(x)w_1(x)$, with $w_0 = w_1 = w$ if utility is state independent.

3 Alternative Models

We look at six different models that capture uncertainty aversion, models which we feel show some promise for evaluation of situations with catastrophic risks. All but the fifth, Chichilnisky’s “sensitivity to rare events,” are built on the Anscombe-Aumann (1963) foundation, implicit in the axioms presented in the Appendix. And all are inspired by the failure of the Savage subjective utility framework to adequately capture behavior in the face of fundamental uncertainty, ambiguity with respect to underlying likelihoods of events. The Ellsberg Paradox, in particular, has provided the challenge to which most of these models respond. And indeed, it is the true uncertainty, ignorance about likelihoods and consequences of truly catastrophic events that makes these models of decision making potentially important to their study.

3.1 ‘Behavioral Probabilities’ and ‘Ambiguity’

3.1.1 Choquet Expected Utility

We begin with the Schmeidler (1989) model of general non-additive valuation of risks using Choquet expected utility (CEU). This captures the idea that an agent’s beliefs are not well specified, hence cannot be represented by a well specified probability. This representation derives from preferences satisfying axioms 1, 2a, 3, 4, and 6 in the Appendix. The representation allows uncertainty about likelihoods, and can reflect aversion to the ambiguity that uncertainty creates. Instead of a probability measure, the state space (S, Σ) is endowed with capacity, $\nu : \Sigma \rightarrow [0, 1]$; $\nu(\emptyset) = 0$, $\nu(S) = 1$; $\forall A, B \in \Sigma, A \subset B \implies \nu(A) \leq \nu(B)$. Each potential act, f , is then evaluated using a

Choquet integral,

$$I_\nu(f) = \int f d\nu \equiv \int_{-\infty}^0 [\nu(f \geq t) - 1] dt + \int_0^\infty \nu(f \geq t) dt, \quad (1)$$

$$f \succsim g \iff I_\nu(f) \geq I_\nu(g)$$

where, in general, $I_\nu(f + g) \neq I_\nu(f) + I_\nu(g)$ for differing acts f, g . A distaste for, an aversion to, ambiguity is captured in the *convexity* (‘supermodularity’) of the capacity ν :

$$\forall A, B \in \Sigma, \nu(A \cup B) + \nu(A \cap B) \geq \nu(A) + \nu(B).$$

An agent with such preferences uses a linear functional, $V(f) \equiv I_\nu(u \circ f) = \int u(f) d\nu$, to evaluate each act/decision that must be taken in the uncertain world. That is, $(f \succsim g) \iff V(f) \geq V(g)$, and an optimal act is any $f^* \in \arg \max_f V(f) = \max_f I_\nu(u \circ f)$.

This formulation can be used to capture other attitudes toward uncertainty/ambiguity, including pure SEU in situations without ambiguity. Indeed, if acts f, g are *co-monotonic*. i.e. $\exists p_\pi$, a probability vector for some permutation of states, π , such that

$$I_\nu(f) = \int_S f dp_\pi = \int_S g dp_\pi = I_\nu(g),$$

then $I_\nu(\cdot)$ becomes additive with respect to f and g ; the comparison involves no ambiguity. Further, it can capture behavioral distortions of known probabilities, such as those revealed in the Allais paradox experiments and modeled in Prospect Theory.

3.1.2 Rank Dependent Expected Utility

A special case of non-additive ‘probability’ can be tractably analyzed using a second model, involving known risks, called Rank Dependent Expected Utility (Quiggin, 1982). In this model, the finite set of outcomes, $x = (x_1, x_2, \dots, x_N)$, is ordered from lowest (worst) to highest (best). Let $p = (p_1, p_2, \dots, p_N)$ be the probabilities of each ‘outcome’, and $F(x_i)$ be the cumulative distribution function of x , evaluated at x_i . The agent is assumed to have a standard Bernoulli utility function over certain outcomes, $u(x)$, but to systematically distort the probabilities that she uses in evaluating prospects/decisions by using a probability weighting function that depends on the value/rank of the outcome. Such a function can capture the systematic deviations found in the experimental literature, such as overweighting or underweighting extreme (i.e. near zero or one) probabilities.

A representation of the preferences of such an agent is given by $V : \mathbb{R}^{2N} \rightarrow \mathbb{R}$, the RDEU function:

$$V(\bar{x}, \bar{p}) = \sum_{i=1}^N u(x_i) h_i(\bar{p}), \quad (2)$$

where $h_i : \mathbb{R}^N \rightarrow [0, 1]$, $i = 1, 2, \dots, N$; $h_i(p) := q[F(x_i)] - q[F(x_{i-1})]$ weights outcome i as a function of the cumulative distribution of x , using a probability weighting function, $q : [0, 1] \rightarrow [0, 1]$, $q(0) = 0$, $q' \geq 0$, that captures the agent’s behavioral understanding of the probabilities. This allows non-linear (as revealed in behavioral experiments) probabilities, while preserving first-order stochastic dominance. Quiggin’s felicitous example is:

$$q(F) = \frac{F^\gamma}{\left(F^\gamma + (1 - F)^{1-\gamma}\right)^{1/\gamma}};$$

$$\gamma \in (0, 1).$$

which ‘overweights’ extreme events [$h_i(\bar{p}) > p_i$, i near 1 or N], generalizing Tversky, Kahneman (1992). Thus it captures empirical behavioral regularity at both the individual and group levels (experimental evidence; Gonzales, Wu, 1999).⁵ Note that a concave $q(F)$ implies classic risk aversion.

The distortion of probabilities that this implies is easy to see graphically. It is clear in the overall shape of Figure 1, where the (dashed) diagonal shows the non-distorted cumulative probabilities. Different levels of distortion are indicated by color: $\gamma = 0.1$ — red; $\gamma = 0.2$ — black; $\gamma = 0.3$ — green; $\gamma = 0.5$ — blue. In Figure 2, we magnify the graph in the vicinity of zero, near the probability of the worst possible outcome, showing that lower γ generates a greater increase in the exaggeration of the perceived probability of the worst event.

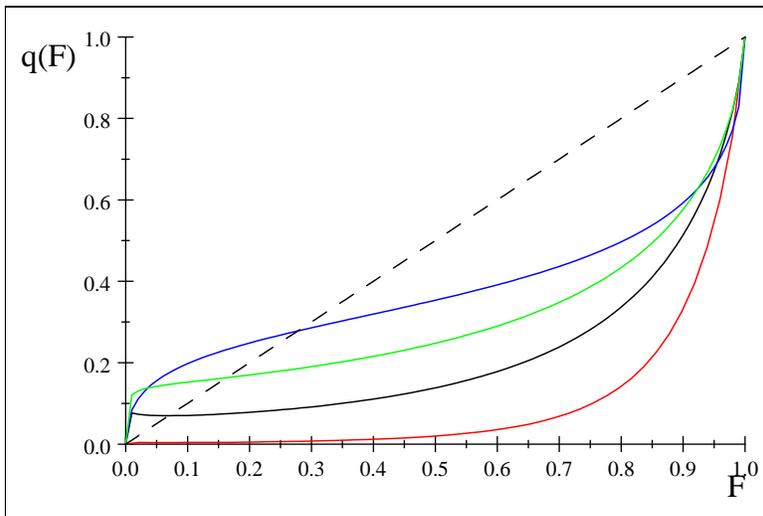


Figure 1: Perception of Outcome Distribution

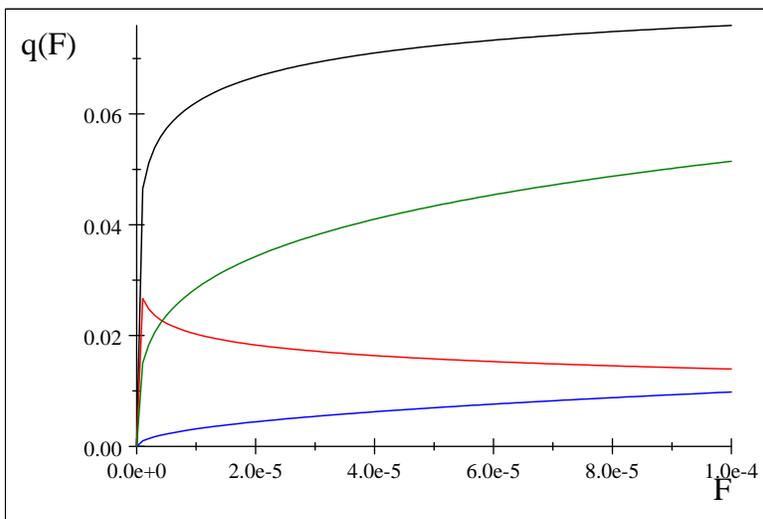


Figure 2: ‘Rare Event’ Probability Perception

⁵Overweighting at the bottom may capture ‘fear’ of the worst outcome, as overweighting at the top may capture ‘hope’ for the best outcome.

Analytic Example. The impact of preferences displaying rank dependent expected utility can be seen in the 2-outcome example. Such an agent dramatically overestimates the likelihood of the ‘catastrophic’ outcome, c , leading to a dramatically diminished valuation (certainty equivalent) of the prospect of facing that catastrophe. Letting $\gamma = 0.2$, $q(0.005) = 0.079453$, $q(0.995) = 0.93456$, and hence $RDE(x) = 11.254$ and $RDEu(x) = 0.777$ with a u -certainty equivalent of 4.489. This is illustrated in Figure 3, where ‘circles’ indicate expected utility and ‘boxes’ RDEU results. Note the non-additivity of the Quiggin probabilities, indicating extra weight at the extremes.

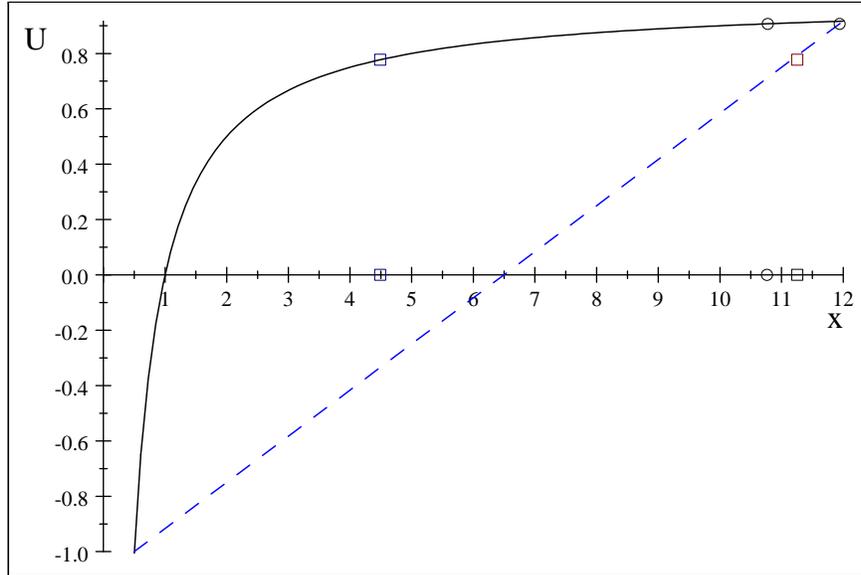


Figure 3: Catastrophic Risk: EU(x) & RDEU(x) Certainty Equivalent

3.1.3 Maximin Expected Utility

Another promising model of preferences in the face of true uncertainty is the maximin expected utility representation of Gilboa and Schmeidler (1989). Rather than ‘distorting’ probabilities, this model captures an agent’s ignorance/uncertainty about the likelihood of events in a set, $\Phi \subset \Delta(\Sigma)$, of potential/conceivable probability distributions over (S, Σ) , with general S , and Φ assumed weak*-compact.⁶ Preferences over outcomes are represented by $u(\cdot)$ — strictly increasing, continuous, weakly concave — and the agent’s evaluation of an ‘act’/decision f , a real-valued, measurable, bounded function, is given by

$$V(f) \equiv \min_{\varphi \in \Phi} \left(\int_S u(f(s)) d\varphi(s) \right). \quad (3)$$

This derives from preferences satisfying axioms 1, 2b, 3, 4, 5, and 6, where axiom 5, a ‘convexity of preferences’ or ‘preference for hedging’ assumption, gives the uncertainty aversion based on supermodularity/convexity of the capacity in CEU. Hence the optimal decision/act is, $f^* \in$

⁶This is the $\sigma(\Delta(\Sigma), B_0(\Sigma))$ topology: a net $\{p_\iota\}_{\iota \in \Upsilon}$ converges to p iff $p_\iota(A) \rightarrow p(A)$, $\forall A \in \Sigma$.

$\arg \max V(f) = \max_f \left\{ \min_{\varphi \in \Phi} \left(\int_S u(f(s)) d\varphi(s) \right) \right\}$, the maximizing act against the minimizing distribution over outcomes.

This formulation provides a cognitive interpretation of Choquet Expected Utility, when $\Phi = \text{core}(\nu) = \{\varphi \in \Delta(\Sigma) \mid \varphi(A) \geq \nu(A), \forall A \subset \Sigma\}$, and $\Delta(\Sigma)$ is the space of all finitely additive probability measures on S . Then probabilities can be understood as based on past experience. However, Φ can be *more general*, containing $\varphi \notin \text{core}(\nu)$ for any capacity ν .

This formulation of preferences clearly separates the agent's 'uncertainty', captured in Φ , and the agent's attitude/aversion towards that uncertainty, captured in $V(\cdot)$. Indeed, it displays a strong uncertainty aversion, an *unwillingness* to place any order, any likelihood, over the distributions in Φ , behaving as if the worst possible distribution there is actually true. The only limits to this pessimism in the face of ambiguity/uncertainty are in the size of the set Φ . Indeed, if Φ is a singleton, this criterion reduces to maximizing (subjective) expected utility — there is no ambiguity. On the other hand, if the agent is entirely ignorant of (or unwilling to contemplate) possible probability distributions, it reduces to the Wald Maximin Criterion:

$$\max_f \left\{ \min_s u(f(s)) \right\}.$$

3.1.4 Smooth Ambiguity Aversion (Second Order Expected Utility)

This representation of preferences in the face of uncertainty/ambiguity extends the prior model by assuming a more sophisticated handling of the possible distributions of outcomes. Rather than assuming the worst, the decision maker places a (subjective) distribution over the set Φ , allowing her to weigh the likelihood of different distributions governing the impact of acts on (payoff relevant) outcomes. One particularly clear model is that of Klibanoff, Marinacci, and Mukerji (KMM, 2005) which parametrizes the agent's uncertainty about a deeper (second-order) 'state', $\theta \in \Theta$, that determines the distribution of states, $\varphi_\theta \in \Delta(\Sigma)$, affecting the payoffs to acts.⁷ The agent's subjective distribution ζ over Θ reflects her uncertainty about the distribution, φ_θ , governing 'events' in S . The attitude toward this 'state risk' is then modeled through introducing another strictly increasing function, $v : \mathbb{R} \rightarrow \mathbb{R}$, which together with $u : X \rightarrow \mathbb{R}$ captures the 'taste for ambiguity' of this agent. As usual, u reflects the agent's 'attitude toward risk'.

Based on the assumption that both first-order and second-order preferences are mutually consistent and have expected utility representations, the decision relevant preference representation over acts, f , becomes:

$$\begin{aligned} U(f) &= v^{-1} \mathbb{E}_\zeta v(u^{-1}(\mathbb{E}_{\varphi_\theta} u \circ f)) \equiv v^{-1} \mathbb{E}_\zeta \phi(\mathbb{E}_{\varphi_\theta} u \circ f) \\ &\equiv v^{-1} \left(\int_\Theta v \left(u^{-1} \left(\int_S u(f) d\varphi_\theta \right) \right) d\zeta(\theta) \right), \end{aligned} \quad (4)$$

where $u^{-1}(\int_S u(f) d\varphi_\theta)$ is the *certainty equivalent* of the gamble induced by decision f .

$$\begin{aligned} (f \succsim g) &\iff U(f) \geq U(g) \\ &\iff v^{-1} \mathbb{E}_\zeta v(u^{-1}(\mathbb{E}_{\varphi_\theta} u \circ f)) \geq v^{-1} \mathbb{E}_\zeta v(u^{-1}(\mathbb{E}_{\varphi_\theta} u \circ g)). \end{aligned}$$

⁷This is derived assuming 'second order' preferences over 'acts' mapping the set of all probabilities over a sufficiently rich 'state space' directly into consequences, that is consistent with first-order preferences over Savage acts from 'states' to consequences, both satisfying the (subjective) expected utility hypothesis. KMM (2005), pp. 1856-9.

‘Acts’ are thus ranked by the ‘certainty equivalent’ (CE) of the induced (by ζ) distribution of the CEs of the ‘lotteries’ in each ‘state’.⁸ This formulation separates ‘ambiguity’ (beliefs) from ‘attitude toward ambiguity’ (tastes): the distribution ζ over Θ , $|\Theta| > 1$, captures ambiguity (beliefs); the composite function $\phi := v \circ u^{-1}$ captures *attitude* toward ambiguity – concavity \iff ‘ambiguity aversion’. As a firmly (subjectively) held belief, ζ also reflects the pessimism/optimism of the agent, with a ‘pessimistic’ agent more heavily weighting the state(s), θ , with the greatest probability of disaster. Here, ambiguity aversion is an aversion to ‘mean preserving spreads’ in the distribution of $\mathbb{E}_{\varphi_\theta} u \circ f$ induced by ζ and f . If, however, the composite function ϕ is linear, then we have ‘ambiguity neutrality’, implying the reducibility of the compound distribution, rendering the representation observationally equivalent to expected utility with the subjective prior ζ .

Analytic Example. The impact of these preferences can again be nicely illustrated in our 2-outcome example of catastrophic risk. Suppose that the agent’s uncertainty about the state of the world is fully captured by 2 possible distributions: in θ_0 the probabilities are (0.05, 0.95) while in state θ_1 they are (0.005, 0.995) as above.⁹ Hence in the worst case (θ_0), the catastrophic state is 10 times as likely as in our base case. Let the agent’s underlying aversion to risk be reflected in the utility function, $u(x)$, as above, and let $\nu(r) := 1 + \frac{x^{1-\gamma}}{1-\gamma} = 1.0 - \frac{1}{3x^3}$ ($\gamma = 4$), reflecting *ambiguity aversion*; $\phi(r) := v \circ u^{-1}(r) = \left(1 - \frac{1}{3\left(\frac{1}{r-1}\right)^3}\right) = 1 - \frac{1}{3}(r-1)^3$ is a clearly concave function. The two ‘valuation’ functions, $u(x)$ over outcomes and $\nu(u^{-1}(\cdot))$ over certainty equivalents of risky prospects, are presented in Figure 4, where the utilities are normalized to be equal in the best of all possible outcomes.¹⁰

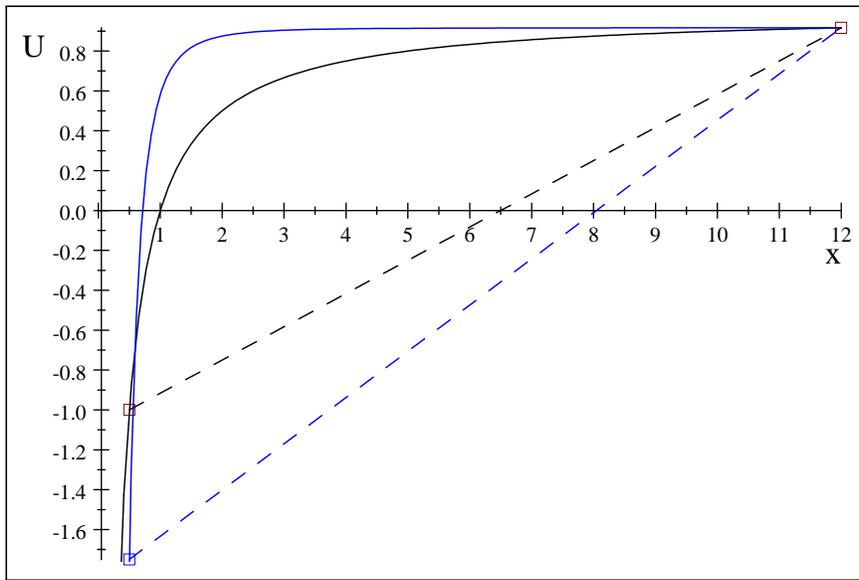


Figure 4: $\nu(x)$ [blue] and $u(x)$ [black], equal at $x = 12$

⁸ $U(f)$ represents identical preferences to $V(f) \equiv \mathbb{E}_\zeta \phi(\mathbb{E}_{\varphi_\theta} u \circ f)$ (KMM, 2005), as v is strictly monotonic. We find it more convenient to work with certainty equivalents in the fixed outcome space than with (subjective) utility values.

⁹Note that this worst case places a lower probability (0.05) on the ‘catastrophe’ than does the RDEU distortion (0.795) above.

¹⁰In this normalization, if the certainty equivalent of a risk were the worst possible outcome, the ‘utility’ would be $\nu(0.5) = -1.75 < u(0.5) = -1$, a true catastrophe. The precise normalization of origin is, however, irrelevant to the decision; only the curvature has meaning.

The circles on $u(x)$ in Figure 5 indicate the expected values and expected utilities of the gambles, and boxes give the certainty equivalents, in ‘states’ θ_0 and θ_1 respectively. The agent’s uncertainty relates to which of these ‘worlds’ obtains, leading her to evaluate this ambiguity by ‘weighting’ these worlds with the distribution ζ .

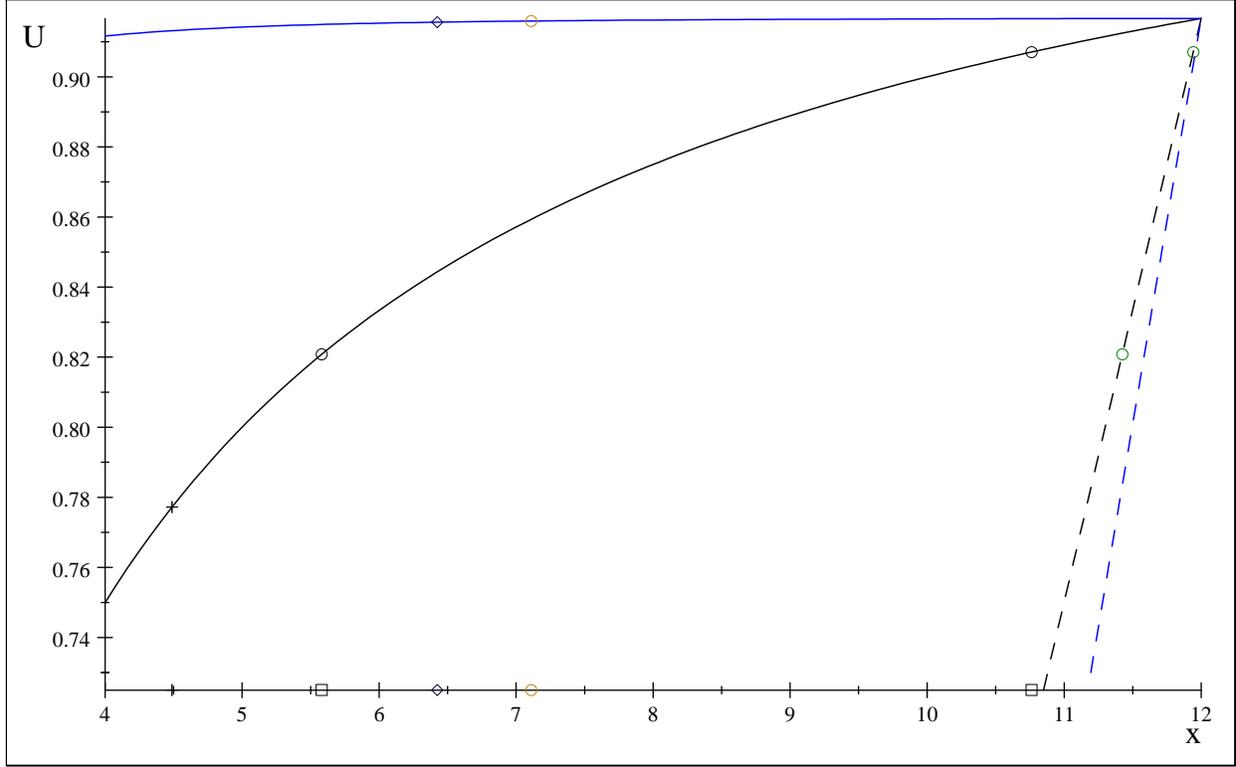


Figure 5: SOEU Evaluations with RDEU Comparison

Here the evaluation function, U , is given by:

$$U \equiv v^{-1} \left(\sum_{\theta} v \left(u^{-1} \left(p_c^{\theta} u(c) + p_n^{\theta} u(n) \right) \right) d\zeta(\theta) \right),$$

that is, it is the certainty equivalent of the ζ -weighting of the ν -values of each u -certainty equivalent in each ‘world’, θ . Let the agent believe $\zeta = (.4, .6)$, i.e. that there is a 40% chance that the truly catastrophic world obtains. Then

$$\begin{aligned} U &= v^{-1} (0.4 \cdot \nu(5.5814) + 0.6 \cdot \nu(10.7623)) \\ &= \nu^{-1} (0.915932711) = 7.110299, \end{aligned}$$

indicated by circles on $\nu(x)$ [the upper curve] and the x-axis in Figure 5. A more pessimistic view, $\zeta = (.6, .4)$, yields $U = 6.42423$, as indicated in the graph with diamonds. Note that both these views are more optimistic than give by the Quiggin probability distortion in the RDEU model; see the + on $u(x)$ in Figure 5.

Indeed, this second order approach remains valid when first-order preferences are assumed to be RDEU (hence non-additive), so that the inner integral becomes a Choquet integral. If we look at the RDEU valuation in the θ_0 world, using $q(F)$ above, we find that the decision maker’s weights assigned to the 2 outcomes sum to only 0.7413, as she assesses $q(0.05) = 0.0702$ and $q(0.95) = 0.6711$. Hence her rank-dependent ‘expectation’ of the outcome is 8.0883 and her RDEU valuation is 0.54498, as illustrated in Figure 6.

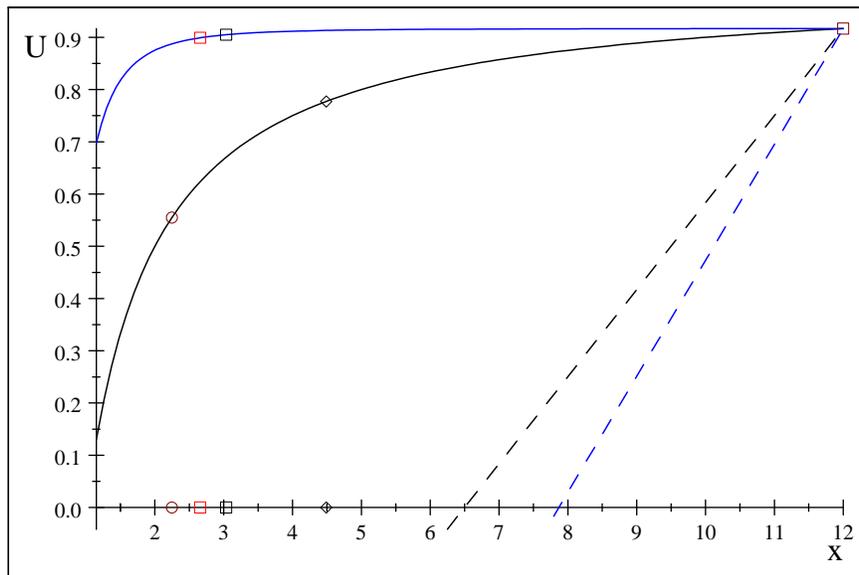


Figure 6: RDEU CEs & SOEU Evaluations

Here we also illustrate the SOEU when $\zeta = (.4, .6)$ — there’s only a 40% chance of the disastrous world in which there is a 5% chance of a catastrophe, rather than just 0.5% — with a ‘box’ on the ‘second order’ utility [upper curve], ν :

$$U = v^{-1}(0.4 \cdot \nu(2.2471) + 0.6 \cdot \nu(10.7623)) = 3.036.$$

With the more pessimistic view, $\zeta = (.6, .4)$, $U = 2.6589$ [lighter box]. Hence with Choquet/Quiggin expected utility over outcomes, the agent places much greater weight on the possibility of disaster, and a much lower evaluation of the world containing such a prospect, than does an agent with SOEU preferences and ζ beliefs.

3.2 “Sensitivity to Rare Events: A Topology of Fear”

This model takes a radically different approach to providing a formal, hence analytically useful, representation for preferences in the face of truly catastrophic risks. Rather than using the Anscombe-Aumann framework as above, this model begins from the von Neuman-Morgenstern axioms of weak preference, independence, monotonicity, continuity and boundedness of utility, and monotone continuity of beliefs (countably additive measure on the space/ σ -field of events). Chichilnisky (2000, 2009, 2010a, 2010b) argues that the resulting countably additive valuation functional, the expected utility representation, cannot, in principle, deal with catastrophic risk, as it necessarily ignores “rare events,” the only events relevant for true catastrophe. Such a functional arises from the assumption of monotone continuity of preferences on the σ -field of events. Continuous bounded utility and monotone continuity imply the impact of any catastrophe, however large, becomes analytically negligible as its likelihood (Lebesgue measure of the event) diminishes to zero. Or, if marginal utility

is unbounded, the potential catastrophic risk swamps all other considerations, however unlikely its occurrence may be (Buchholz, Schymura (2010)). Hence, she argues, to analyze truly catastrophic risks, we need to break free of the standard model, and develop new analytic foundations for that analysis.

To do so, Chichilnisky introduces axioms requiring “equal treatment” of both ‘rare’ and ‘frequent’ events, and dispenses with monotone continuity as contradicting sensitivity to rare events. This requires a richer space of ‘states’ than other formulations — both unboundedness and cardinality no less than the continuum — one rich enough to capture all possible consequences, however unlikely, of any ‘act’. To develop a basic representation, the simplest such space, $S = \mathbb{R}^1$, is used. It is endowed with Lebesgue measure $\mu(s)$, hence assuming that the underlying ‘states’ are ‘equiprobable’, with the likelihood/probability of an event depending on the ‘number’ of states in it. Functions $f, g : \mathbb{R} \rightarrow \mathbb{R}$ are ‘acts’ generating “lotteries,” giving the ‘utility’ payoffs $u(f(s))$ and $u(g(s))$ to each ‘state’ s . The consequences, $u \circ f$, are assumed to be a.s. (λ) bounded functions. To insure that the impact of both frequent ($\mu(E) > 0$) and rare ($E \subset S, \mu(S) = 0$) events are captured, the space of ‘acts’, $\mathcal{F} : = L_\infty(\mathbb{R})$ [essentially bounded functions], is endowed with the topology generated by the esssup norm, $\|\cdot\|_\infty$, felicitously named the “topology of fear” (Chichilnisky, 2009).

To generate a preference representation, Chichilnisky (2009, 809-810) applies a new axiom in the classic decision theory framework (Arrow, 1971), “sensitivity to rare events.” This gives 3 (summary) axioms for a ranking (preference representation):

Axiom 1. The ranking $W : L_\infty(\mathbb{R}) \rightarrow \mathbb{R}$ is linear and $\|\cdot\|_\infty$ -continuous;

Axiom 2. The ranking $W : L_\infty(\mathbb{R}) \rightarrow \mathbb{R}$ is sensitive to rare events: the ordering of f and g can depend on consequences in events of arbitrarily small (Lebesgue) measure;¹¹

Axiom 3. The ranking $W : L_\infty(\mathbb{R}) \rightarrow \mathbb{R}$ is sensitive to frequent events: the ordering of f and g is dependent on payoffs in events of sufficiently large (Lebesgue) measure;¹²

Under these axioms, a preference representation for the evaluation of acts with potential extreme consequences is shown to exist.¹³ The representation is a bounded linear functional on $L_\infty(\mathbb{R})$,

$$W(f) = \lambda \int_{\mathbb{R}} u(f(s)) \phi_1(s) ds + (1 - \lambda) \langle u \circ f, \phi_2 \rangle, \quad (5)$$

where $\lambda \in (0, 1]$, $\phi_1, \phi_2 \in L_\infty^*$ — continuous linear functionals on L_∞ ; $\phi_1 \in L_1(\mathbb{R})$, $\int_{\mathbb{R}} \phi_1(x) dx = 1$, and ϕ_2 is a purely finitely additive functional. The parameter λ is an essential aspect of the underlying preferences:

$$\begin{aligned} (f \succeq g) &\iff \lambda \int_{\mathbb{R}} u(f(s)) \phi_1(s) ds + (1 - \lambda) \langle u \circ f, \phi_2 \rangle \\ &\geq \lambda \int_{\mathbb{R}} u(g(s)) \phi_1(s) ds + (1 - \lambda) \langle u \circ g, \phi_2 \rangle. \end{aligned}$$

λ reflects the decision maker’s ‘belief’ in the reality of extremely rare events, while ϕ_2 provides her ‘evaluation’ of the consequences in such events. Both arise from a ‘sensitivity’ to rare events, with

¹¹Formally, the negation of “insensitivity to rare events:” $\forall f, g, \exists \epsilon = \epsilon(f, g) > 0$, s.t. $W(f) > W(g) \iff W(f') > W(g') \forall f', g'$ satisfying $f' = f, g' = g$ a.e. on $A \subset \mathbb{R}$ when $\mu(A^c) < \epsilon$.

¹²Formally, the negation of “insensitivity to frequent events:” $\forall f, g, \exists \delta = \delta(f, g) \in (0, 1)$, s.t. $W(f) > W(g) \iff W(f') > W(g') \forall f', g'$ satisfying $f' = f, g' = g$ a.e. on $A \subset \mathbb{R}$ when $\mu(A^c) > \delta$.

¹³The proof is non-constructive, using the axiom of choice.

ϕ_2 placing a value on unmeasurable events within those of Lebesgue measure zero, and λ providing the ‘utility weight’ placed on the “normal” outcomes of the ‘lottery’ f . Notice that, in the absence of ‘extreme outcomes’, $W(f)$ reduces to classic expected utility. This could also occur if the agent refuses to consider the possibility of extreme events, or is “paralyzed by fear,” hence setting $\lambda = 1$.

This model would seem to provide a very direct way to deal with true catastrophe, but its mathematical complexity, the non-constructive derivation, and general lack of explicit analytic form for the purely finitely additive (pfa) valuation functional make it hard to see how this representation can be applied, aside from elucidating some general principles.¹⁴ Thus Chichilnisky has proposed a finite-state version of the representation, which however lacks the same firm mathematical foundation.

3.2.1 Finite-State Version with Extreme Event

Let $S, X \subset \mathbb{R}^S$, and \mathcal{F} be as in the ‘ambiguity’ models above, and $s^* \in S$ be a ‘rare’ catastrophic state if its probability $\phi_{s^*} < \varepsilon$ for some arbitrarily small $\varepsilon > 0$. By analogy, we have the preference representation, $W : \mathbb{R}^S \rightarrow \mathbb{R}$,

$$W(f) = \lambda \cdot \langle \phi, f \rangle + (1 - \lambda) \cdot \min_s f_s, \quad (6)$$

which puts extra ‘weight’ on the catastrophic outcome, $\min_s f(s)$. Here the first term is an ‘expected utility’ based on the ‘subjective’ probability vector, ϕ , and the utility vector, $u(f)$, resulting from the act, while the second puts independent weight on the worst that can happen. Maximizing $W(f)$ trades off maximization of expected utility against the minimization of catastrophic loss. Again, λ is a critical preference parameter, which might be derived from a constrained optimization formalizing that trade-off:

$$\max_{f \in F} \langle \phi, f \rangle \quad s.t. \quad \min_s f_s \geq \bar{u}, \quad (7)$$

or

$$\max_{f \in F} \min_{\mu \geq 0} \left[\langle \phi, f \rangle + \mu \left(\min_s f_s - \bar{u} \right) \right] \equiv \max_{f \in F} \min_{\mu \geq 0} L(f, \mu),$$

giving $\lambda = \frac{1}{1+\mu}$.

This formulation satisfies analogies of Axioms 2 and 3 above, but not Axiom 1: it is not continuous (Chichilnisky, 2010), raising a question about its necessity and uniqueness. It also must assume (for relevance) that s^* at which $\min_s f_s$ occurs (the ‘catastrophe’) is ‘rare’: $\varepsilon > \phi_{s^*} > 0$. However, it provides a wide range of valuations of potentially catastrophic lotteries, from standard EU ($\lambda = 1$) to the Wald maximin criterion ($\lambda = 0$), depending on the (subjective) value of λ . Finally, it is easier to apply than the rich model, while maintaining its intuition, and it appears to be justified by the axioms supporting the final model of preferences that we explore below, “variational preferences.”

Analytic Example. The implications of this formulation can be easily seen in our 2-state example. There the probabilities and outcomes associated with the some act are given, so we parametrize the preference representation by λ :

$$\begin{aligned} W(f; \lambda) &= \lambda \cdot \langle \phi, f \rangle + (1 - \lambda) \cdot \min_s f_s \\ &= \lambda (0.005u(c) + 0.995u(n)) + (1 - \lambda) u(c). \end{aligned}$$

¹⁴Some pfa functionals can be expressed as limits, as in the example Chichilnisky (2009, p. 814) gives. One satisfactory infinite dimensional application has been made in Figuières, Tidball (2012) where limiting outcomes are evaluated. In the case of pure uncertainty, and it might be possible to isolate the valuation of ‘rare’ events along an ultrafilter net, but I do not see clearly how to do that.

Hence $W(f; \lambda) = 1.90708\lambda - 1.0$, $W(f; 1) = 0.90708$, and $W(f; 0) = -1.0$, showing the vast range of valuations that variation of λ generates. When $\lambda = 0.5$, $W = -0.04646$, and the certainty equivalent of this lottery becomes 0.9556, clearly reflecting “fear” of the ‘bad’ outcome. It is depicted in Figure 7, where boxes show the expected utility, the certainty equivalent, and its utility. Indeed, the appropriate choice of λ can replicate the evaluation of any of our preference representations. For example, $\lambda = 0.92497 \implies W = 0.764$, which is identical to the *RDEU* (with $\gamma = 0.2$) evaluation, reflecting just a little “fear” — see Figures 3 and 7.

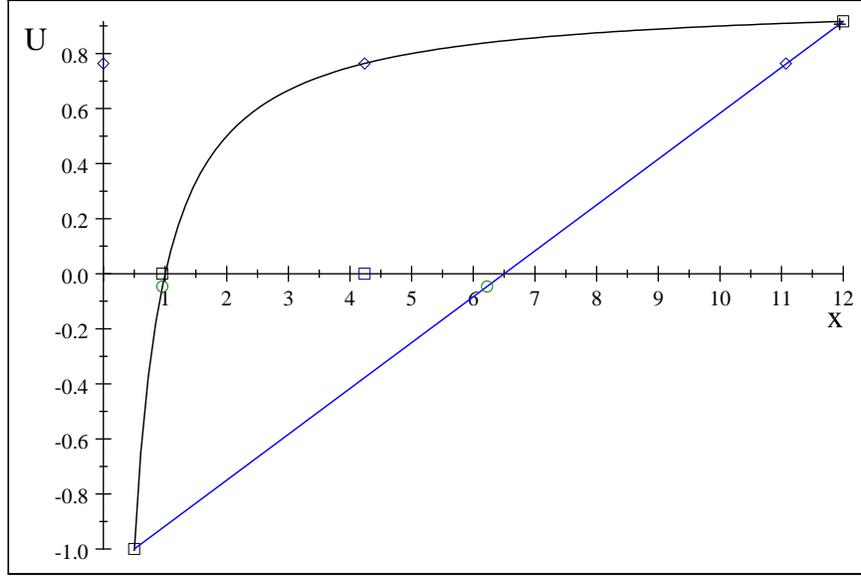


Figure 7: ‘Topology of Fear’ & RDEU(x) Certainty Equivalents

Clearly this representation can generate more extreme responses to extreme hazards than the other modifications of expected utility.

3.3 Variational Preferences : An Umbrella Model?

These preferences were introduced by Maccheroni, Marinacci, and Rustichini (2006) as a unifying framework for understanding the various models of decisions in the face of ambiguity based on a common behavioral (axiomatic) foundation. The models encompassed include those above, and the “multiplier preferences” model of Hansen and Sargent (2001), which is motivated by model uncertainty in control problems. These are a ‘class’ of preferences that can be specialized to cover most other forms of ambiguity respecting preferences, as discussed in Strzalecki (2011), as well as expected utility when the agent is ‘ambiguity neutral’. In addition, variational preferences support the representation suggested for the finite state Chichilnisky model (6).

The common behavioral foundation is laid out in the notation and axioms in the Appendix. Under Axioms 1 through 6, preferences \succsim defined on \mathcal{F} are representable by

$$V(f) = \min_{p \in \Delta(\Sigma)} \left[\int u(f) dp + c(p) \right], \quad (8)$$

where $u : X \rightarrow \mathbb{R}$ is an affine utility; $c(p) : \Delta \rightarrow [0, 1]$ is an ‘index’ of ambiguity aversion, a convex lower-semicontinuous function with $\inf_{\Delta} c(p) = 0$. Over the space of countably additive probabilities, $\Delta^{\sigma}(\Sigma)$, $c(p) = \sup_{f \in \mathcal{F}} (u(x_f) - \int u(f) dp)$, where x_f is the certainty equivalent of

act f , and, with Axiom 7, is unique. $c(p)$ is a “penalty” on less likely distributions, so lower $c(p)$ reflects higher ‘ambiguity aversion’.

The other preference representations can be generated by specifying the form of the ‘ambiguity index’, $c(p)$. For example, multiple priors (maximin) preferences result from $c(p) = \delta_C(p) = \begin{cases} 0, & \text{if } p \in C \\ \infty, & \text{otherwise} \end{cases}$, “multiplier preferences from $c(p) = \theta R(p \parallel q)$ where $R(p \parallel q)$ is the relative entropy of p with respect to $q \in \Delta(\Sigma)$, a fixed, countably additive, non-atomic measure, and $c(p) = \theta G(p \parallel q)$ gives mean-variance preferences, with $V(f) = \int f dq - \frac{1}{2\theta} \text{Var}(f) = \min_{p \in \Delta(\Sigma)} \int f dp + \theta G(p \parallel q)$, where $G(\cdot \parallel q) : \Delta \rightarrow [0, 1]$ is the relative Gini concentration index.¹⁵ All other variational preferences are ambiguity averse, as Axiom 5 indicates a “preference for hedging.”

Defining \succsim_1 as “more ambiguity averse” than \succsim_2 if, $\forall f \in \mathcal{F}$ and $x \in X$, $f \succsim_1 x \implies f \succsim_2 x$, it is clear that $\{c_1 \leq c_2 \text{ for affine equivalent } u_i, i = 1, 2\} \Leftrightarrow \{\succsim_1 \text{ is more ambiguity averse than } \succsim_2\}$. “Ambiguity neutral” (subjective) expected utility, with prior q , arises from $c(p) = \delta_q(p) = \begin{cases} 0, & \text{if } p = q \\ \infty, & \text{otherwise} \end{cases}$; this is the model with minimal ambiguity aversion. the case of maximal ambiguity aversion in this model is given by the Wald criterion: $c_m(p) = 0, \forall p \in \Delta$, so that

$$\begin{aligned} f \succsim g &\Leftrightarrow \min_{p \in \Delta(\Sigma)} \int u(f) dp \geq \min_{p \in \Delta(\Sigma)} \int u(g) dp; \\ f \succ g &\Leftrightarrow \min_{s \in S} u(f(s)) \geq \min_{s \in S} u(g(s)). \end{aligned}$$

Other cases of less extreme ambiguity aversion can be defined using a convex combination of these extremes. Let $\alpha \in (0, 1)$, and

$$V(f) = (1 - \alpha) \int u(f) dq + \alpha \min_{s \in S} u(f(s)).$$

Here, if $c(p) = \min_{p_1, p_2 \in \Delta} \{(1 - \alpha) c_q(p_2) + \alpha c_m(p_1) : (1 - \alpha) p_2 + \alpha p_1 = p\} = \delta_{(1-\alpha)q + \alpha\Delta}(p)$, then this is a special case of variational preferences that gives the same representation as the finite-state “topology of fear” preferences.

Analytic Example. This model can clearly replicate the outcomes of any of the above representations in our 2-outcome example; it can replicate any degree of ‘aversion to uncertainty’ in the face of a potentially catastrophic event. So we turn to the more elaborate (3-outcome) example with distinct outcomes in each of the two states, normal ($s = 1$) and catastrophic ($s = 0$), set up above (Section 2.1.2). Let there be 3 distinct acts mapping S into $\Delta(X)$, $X = \{0.5, 3, 9\}$:

$$\mathcal{F} = \left\{ \begin{array}{ll} f : & f_0 = (.25, .75, 0) \quad f_1 = (0, .8, .2) \\ g : & g_0 = (.5, .5, 0) \quad g_1 = (0, .6, .4) \\ h : & h_0 = (.7, .3, 0) \quad h_1 = (0, .5, .5) \end{array} \right\}. \quad (9)$$

These acts differ in mitigation effort with respect to the potential disaster, $x = 0.5$. Act f takes the possibility of disaster most seriously, placing the greatest effort on mitigation, diverting resources to that effort, and hence reducing the likelihood of the ‘best’ configuration of outcomes. Act h pays little attention to the possibility of a disastrous outcome, ‘maximizing’ the likelihood of the ‘best’ outcomes, as well as of the ‘catastrophe’; it might be considered “business as usual” behavior. Finally, g cautiously compromises between these two approaches.

¹⁵Maccheroni, Marinacci, and Rustichini (2006), pp. 1449-50.

To illustrate the implications of variational preferences for choice in the face of ‘catastrophe’, we adapt a variational preferences representation with state dependent preferences:

$$V(a) = \min_{p \in [0,1]} \{pE_0 \{u_0(a_0)\} + (1-p)E_1 \{u_1(a_1)\} + c(p)\}, a = f, g, h,$$

and let

$$c(p) = 5(p - 0.2)^2, p \in [0, 0.6]; c(p) = M, p \in (0.6, 1].$$

This assumes the agent believes that the ‘most realistic’ estimate the probability of the catastrophic state is $p = 0.2$, and that $p > 0.6$ is ‘impossible’. We also assume that the agent is more risk averse in the ‘catastrophic’ state, and consider 2 cases: (i) Utility level independent of ‘state’ in the most likely configuration of outcomes, $x = 3$; (ii) Utility level is also “state dependent” even when different ‘states’ yield same outcomes ($x = 3$):

- (i) $u_0(x) = 1.15416 + \frac{x^{1-\gamma}}{1-\gamma}; u_1(x) = \ln(x); \gamma = 3; u_0(3) = u_1(3);$
(ii) $u_0(x) = 0.5 + \frac{x^{1-\gamma}}{1-\gamma}; u_1(x) = \ln(x); \gamma = 3; u_0(3) < u_1(3).$

Note that the coefficient of relative risk aversion is 3 in state $s = 0$ and 1 in state $s = 1$. The consequences of the 3 acts, f, g, h , (diamond, box, and circle, respectively, in the graphs) in the catastrophic state, $s = 0$, are illustrated in Figures 8 and 9 in terms of expected utilities (and values) and certainty equivalents, with f yielding the highest value.

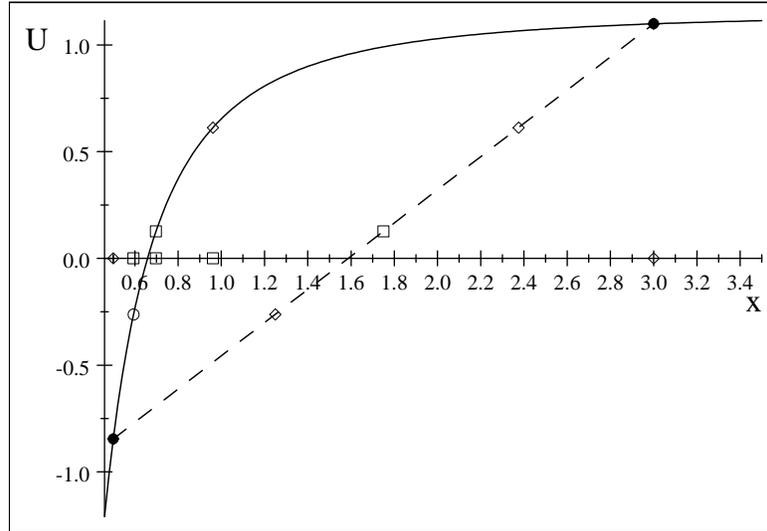


Figure 8: ‘Catastrophe State’ lotteries, $u_0(3) = u_1(3)$

Note that the certainty equivalents remain the same despite the distinctly lower utility levels in case (ii), showing the irrelevance of *that* normalization *within a state*. But that irrelevance vanishes as soon as we study the full consequences, evaluating the outcomes of various acts across states. Case (i) is illustrated in Figure 9, where the green lines connect the outcomes each of the acts f, g, h (least to highest spread) generates in each of the states.

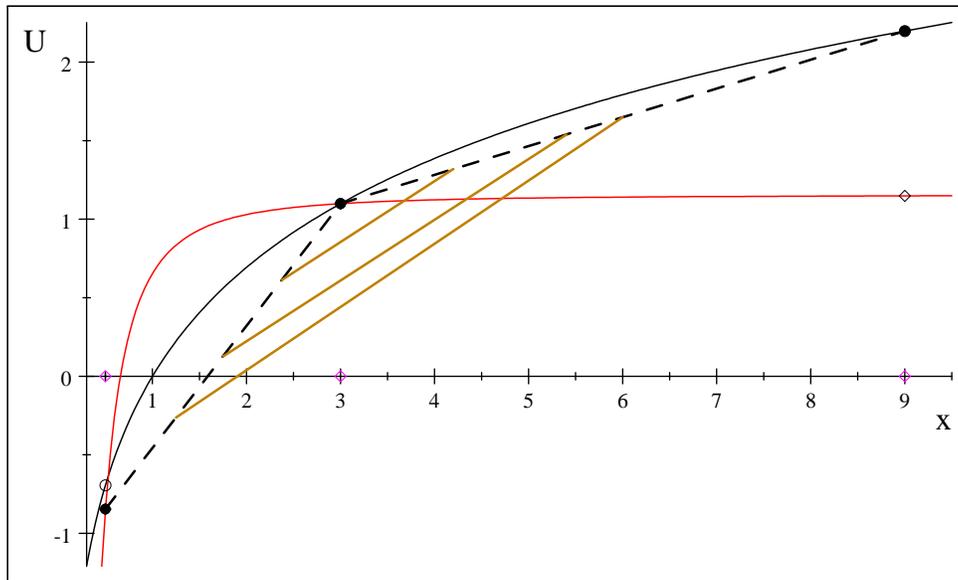


Figure 9: $u(3)$ normalization, and the span of ‘act’ lotteries.

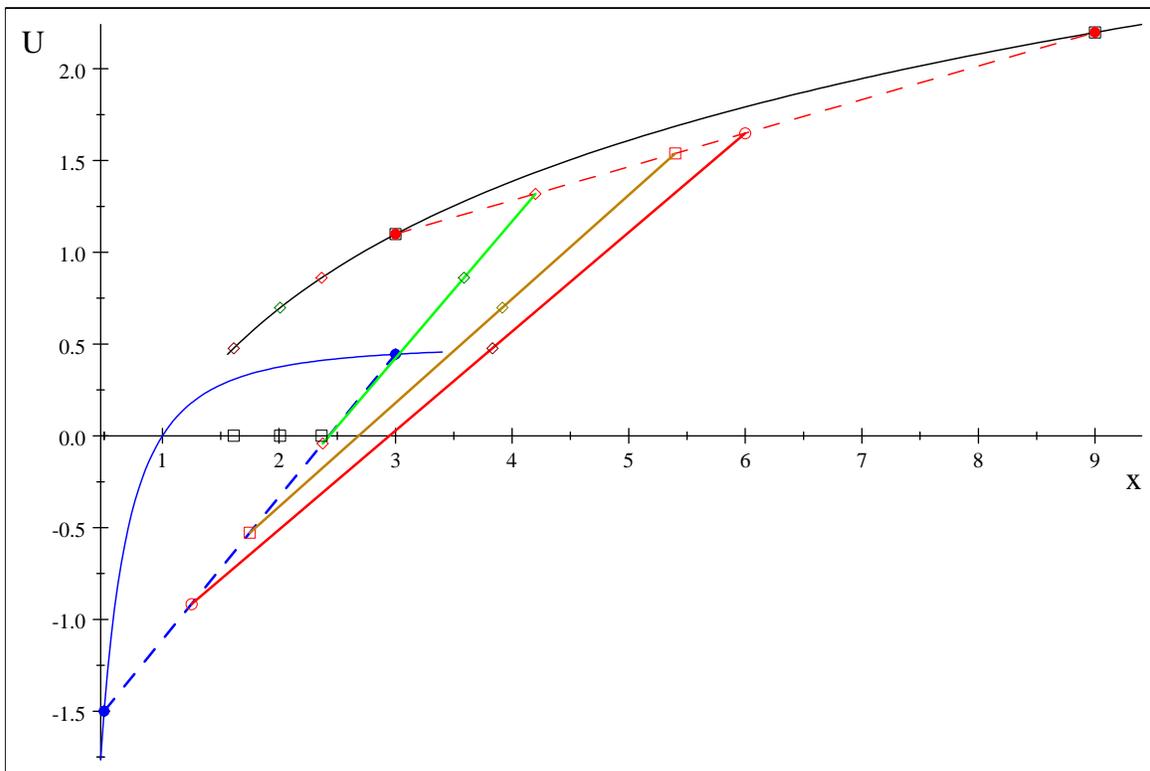


Figure 10: $u_0(x) < u_1(x)$ – expected values, CEs and utilities of f, g, h .

When we take the loss of utility/wellbeing seriously in the catastrophic state, $u_0(x) < u_1(x)$ and more risk averse, the ‘stakes’ in the decisions/acts become much more serious, as can be seen in Figure 10. Here the expected values of outcomes, as a function of p , of the acts f, g , and h are

illustrated, where

$$\begin{aligned} V(f) &= 0.95385 & \text{at } p &= .33600; & \text{Certainty Equivalent: } & 2.3664; \\ V(g) &= 0.91151 & \text{at } p &= .40658; & \text{Certainty Equivalent: } & 2.00999; \\ V(h) &= 0.80615 & \text{at } p &= .45646; & \text{Certainty Equivalent: } & 1.611532, \end{aligned}$$

and $f \succ g \succ h$. Hence ambiguity aversion has led this agent to choose the most cautious, least remunerative act, f . With complete faith in her ‘best estimate’, reflecting absence of ambiguity, $p = 0.2$ [$\Leftrightarrow c(p) = 0$], and $h \succ g \succ f$, as $V(f) = 1.0463$, $V(g) = 1.1249$, and $V(h) = 1.1350$. This leads her to choose the least cautious, most remunerative, act, to pursue “business as usual.” If the agent’s preferences are described by the MEU (maximin) representation, then she believes $p = 0.6$, and $f \succ g \succ h$ as $V(f) = 1.3023$, $V(g) = 1.0986$, and $V(h) = 0.90917$.

It is straightforward to adjust this representation to incorporate RDEU probability distortions or “smooth ambiguity” preferences working with the certainty equivalents of the underlying ‘act lotteries’ and their ‘ambiguity weightings’. More interesting is the impact of the finite version of the “topology of fear” model (Chichilnisky, 2010b). As noted above, this arises from the representation (8)¹⁶, $V(f) = \min_{p \in \Delta(\Sigma)} \int u(f) dp + c(p)$, with

$$c(p) = \min_{p_1, p_2 \in \Delta} \{ \lambda c_q(p_2) + (1 - \lambda) c_m(p_1) : \lambda p_2 + (1 - \lambda) p_1 = p \} = \delta_{\lambda q + (1 - \lambda)\Delta}(p),$$

where $\delta_q(p) = \begin{cases} 0, & \text{if } p = q \\ \infty, & \text{otherwise} \end{cases}$, giving for any act a ,

$$\begin{aligned} V(a) &= \lambda \int u(a) dq + (1 - \lambda) \min_{s \in S} u(a(s)). \\ &= [q E_0 \{u_0(a_0)\} + (1 - q) E_1 \{u_1(a_1)\}] + (1 - \lambda) E_0 \{u_0(a_0)\}. \end{aligned} \tag{10}$$

In our example, we need to specify the agent’s (subjective) distribution over S , q , and the relevant expected utility in each state for each act. Using the ‘best guess’ from above, $q(s_0) = 0.2$, and the lottery generated by each act (9), we can calculate for each act the expected utility in each state, the impact of λ on the overall evaluation, $V(\cdot; \lambda)$, and the evaluation and certainty equivalent for $\lambda = .5$.

s	$Eu(f)$	$Eu(g)$	$Eu(h)$
0	-0.0417	-0.5278	-0.9167
1	1.3183	1.5381	1.6479
$V(\cdot; \lambda)$	$1.088\lambda - 0.0417$	$1.6527\lambda - 0.5278$	$2.0517\lambda - 0.9167$
$V(\cdot; .5)$.5023	.29855	.10915
Cert Equiv	1.6525	1.3479	1.1153

Clearly, if the agent puts full utility weight on the impact of the catastrophe ($\lambda = 0$), the most cautious act f is optimal, while if the agent places no special weight on the catastrophic outcome ($\lambda = 1$), the ordering is reversed, $h \succ g \succ f$, the least cautious being the best. That remains the ordering as long as $\lambda > .9747$ [the agent feels little fear], below which [$\lambda \in (.90796, .9747)$] $g \succ h \succ f$. For λ below 0.8608, the most cautious strategy again becomes dominant: $f \succ g \succ h$. Thus only if the agent puts little weight (less than 0.14) on the catastrophic state will she opt for any but the most cautious act. Finally note that the certainty equivalents of these acts for $\lambda = 0.5$ are given in terms of the ‘normal state’ utility in the final row of the table above. In all cases they are substantially below those in our variational preferences example depicted in Figures 9 and 11, indicating the extreme ambiguity aversion of preferences incorporating fear of a catastrophe.

¹⁶Letting $\lambda = 1 - \alpha$ for notational consistency with that model.

4 Lessons and Conclusions

In this note we have reviewed and illustrated six different models of how decision makers might evaluate catastrophic risks, how their preferences deal with uncertain prospects. These models are grounded in a set of similar behavioral axioms, with differences in axioms identifying them. Each model shares a behavioral foundation based on a distaste for (‘fear of’) uncertainty/ambiguity, i.e. a lack of even probabilistic knowledge of, or firm beliefs about, the underlying events driving (potential) outcomes. And all yield a greater degree of caution, a desire to “hedge” against the unknown, than does the model of a well-informed agent with (subjective) expected utility preferences. Hence they each provide a cognitive explanation for the predictive shortcomings of the SEU model, providing a rational basis for observed human behavior in the face of true uncertainty.

In doing so, these models would appear better able to provide predictions of the human responses to potential catastrophe than the SEU model does. In an economic analysis of catastrophic risk, where the decisions, actions, and reactions of millions of (potentially impacted) individuals must be considered in policy formation, such models should be a critical tool. Whatever the policy considered, risk mitigation, management or reaction planning, an equilibrium analysis of individual behavioral responses which may further, undercut, or even negate policy measures, is essential. These models, and in particular that of “variational preferences,” could be extremely helpful in that analysis.

Indeed, the cognitive structures that these models reflect may also be relevant to decision makers responsible for developing and implementing social policy with regard to catastrophic risks. On the other hand, it might be argued that relevant social preferences should differ from individual, captured in the models above, much as social discount rates should differ from those of individuals. That is an argument and analysis that goes beyond the present discussion. Even if the optimal rational social choice criterion is best reflected in minimally (or zero?) discounted SEU, where likelihoods (the ‘state’ distribution) are based on the best scientific evidence, each of these models admits SEU as an “ambiguity neutral” boundary case. Further, they each indicate the direction of potential bias that can arise from “uneasiness” about that scientific evidence/knowledge. With that in mind, we feel that the final model presented, that of “variational preferences,” provides the most practical, analytically tractable model representation of ambiguity/uncertainty respecting preferences. Through its flexible specification of a convex ‘ambiguity penalty’ function it captures most of the other representations.

These representations with their clear axiomatic foundations lend themselves to experimental testing or which model better captures agent behavior in the face of very unknown probability, high cost gambles. It is now also important to develop specific models of situations involving great uncertainty and (potentially) vast costs to the decision maker, and to bring such models to data. One example might be the analysis of the mixed response in threatened populations to (mandatory) evacuation orders.¹⁷ Finally, work is needed placing these models of individual behavior within an equilibrium framework within which policy decisions must be made. All these remain on our research agenda.

5 APPENDIX: Axiomatic Foundations of Representations

5.1 Notation

The following notation is common to all the models presented.

¹⁷See, for example, the studies of Dow and Cutter (2000) and Lindell, Lu, and Prater (2005).

(S, Σ)	– states, with <i>algebra</i> of events;
$\Delta(\Sigma)$	– all finitely additive probability measures on (S, Σ)
X	– set of consequences, convex $\subset \mathbb{R}^n$
\mathcal{F}	– set of ‘acts’ $f : S \rightarrow X$, simple functions
$B_0(\Sigma)$	– real valued, measurable simple functions
$B(\Sigma)$	– supnorm closure of $B_0(\Sigma)$
$B_0(\Sigma, K)$	– simple maps into $K \subset \mathbb{R}$
$u : X \rightarrow \mathbb{R}$	– affine utility
$I : B_0(\Sigma, u(X)) \rightarrow \mathbb{R}$	– “certainty equivalent” representing ‘beliefs’

In the examples, we simplify to:

$S = \{0, 1\}$,	the ‘catastrophic’ and ‘normal’ states, respectively;
$X = X_0 \cup X_1 \subset \mathbb{R}$;	$X_0 \cap X_1 \neq \emptyset$;
$\Delta(X)$	a set of finitely additive measures on X ;
$\mathcal{F} =$	$\{f : S \rightarrow \Delta(X) \mid f(0) \in \Delta X_0, f(1) \in \Delta X_1\}$
$u_s : X_s \rightarrow \mathbb{R}$,	a ‘state-dependent’ utility function.

5.2 Axioms

1. Weak Order: \succsim complete, transitive on \mathcal{F} .
2. Weak Certainty Independence: $f, g \in \mathcal{F}$, $x, y \in X$, $\alpha \in (0, 1)$

$$\alpha f + (1 - \alpha)x \succsim \alpha g + (1 - \alpha)x \implies \alpha f + (1 - \alpha)y \succsim \alpha g + (1 - \alpha)y.$$

- (a) Independence: $f, g, h \in \mathcal{F}$, $\alpha \in (0, 1)$

$$f \succsim g \Leftrightarrow \alpha f + (1 - \alpha)h \succsim \alpha g + (1 - \alpha)h.$$

- (b) Comonotonicity Independence (Schmeidler, 1989): $f, g, h \in \mathcal{F}$ pairwise *comonotonic*, $\alpha \in (0, 1)$

$$f \succ g \Leftrightarrow \alpha f + (1 - \alpha)h \succ \alpha g + (1 - \alpha)h;$$

- f and g are *comonotonic* if $\neg \exists s, t \in S, f(s) \succ g(s) \wedge g(t) \succ f(t)$.

- (c) Certainty Independence (CI: Gilboa-Schmeidler, 1989):

$$f \succsim g \Leftrightarrow \alpha f + (1 - \alpha)x \succsim \alpha g + (1 - \alpha)x.$$

Lemma 1 \succsim satisfies CI iff $f, g \in \mathcal{F}$, $x, y \in X$, $\alpha, \beta \in (0, 1)$

$$\alpha f + (1 - \alpha)x \succsim \alpha g + (1 - \alpha)x \implies \beta f + (1 - \beta)y \succsim \beta g + (1 - \beta)y.$$

3. Continuity: $\forall f, g, h \in \mathcal{F}$, $\{\alpha \in [0, 1] \mid \alpha f + (1 - \alpha)g \succsim h\}$ and $\{\alpha \in [0, 1] \mid h \succsim \alpha f + (1 - \alpha)g\}$ are closed.
4. Monotonicity [state independence condition]: $f, g \in \mathcal{F}$, $f(s) \succsim g(s)$, $\forall s \implies f \succsim g$.
5. Uncertainty/Ambiguity Aversion: $f, g \in \mathcal{F}$, $\alpha \in (0, 1)$

$$f \sim g \implies \alpha f + (1 - \alpha)g \succsim f.$$

6. Nondegeneracy: $f \succ g$ for some $f, g \in \mathcal{F}(\Delta(Z))$.

7. Unboundedness: $\exists x \succ y \in X$ s.t. $\forall \alpha \in (0, 1)$, $\exists z \in X$ satisfying either (i) $y \succ \alpha z + (1 - \alpha)x$ or (ii) $\alpha z + (1 - \alpha)y \succ x$.
8. Monotone Continuity: If $x \in X$ [$\pi \in \Delta(Z)$] $\{E_n\}_{n \geq 1} \in \Sigma$ with $E_1 \supseteq E_2 \supseteq \dots$ and $\bigcap_{n \geq 1} E_n = \emptyset$, then $f \succ g \implies \exists n_0 \geq 1$ s.t. $x E_{n_0} f \succ g$.
 - Axioms 1, 3, 4 imply $\forall f \in \mathcal{F}$, $\exists x_f \in X$, a “certainty equivalent.”
 - Ambiguity Neutrality \iff SEU: $U(f) = \int u(f) dp$.

References

- [1] F.J. Anscombe, R.J. Aumann, “A definition of Subjective Probability,” *Annals of Mathematical Statistics*, **34**, 199-205, 1963.
- [2] W. Buchholz, M. Schymura (2010), “Expected Utility Theory and the Tyranny of Catastrophic Risks,” 25 August 2010.
- [3] G. Chichilnisky (2000), “An Axiomatic Approach to Choice Under Uncertainty with Catastrophic Risks,” *Resource and Energy Economics*, **22**(3), 221-231.
- [4] G. Chichilnisky (2009), “The Topology of Fear,” *J. Math Econ*, **45**, 807-816, 2009.
- [5] G. Chichilnisky (2010a), “Subjective Probability with Black Swans,” *J. Prob. & Stats*, 2010.
- [6] G. Chichilnisky (2010b), “Catastrophic Risks with Finite or Infinite states,” *International Journal of Ecological Economics and Statistics*, **23**(11F), 3-18, 2010.
- [7] G. Chichilnisky, P. Eisenberger (2010), “Asteroids: Assessing the Catastrophic Risks,” *J. Prob. & Stats*, Article ID 954750, 2010
- [8] K. Dow, S.L. Cutter (2000), “Public Orders and Personal Opinions: Household Strategies for Hurricane Risk Reduction,” *Global Environmental Change, Part B*, **2**(4), 143-155, (2000).
- [9] C., Figuières, M. Tidball (2012), “Sustainable Exploitation of a Natural Resource: a satisfying use of Chichilnisky’s criterion,” *Economic Theory*, **49**, 243-265.
- [10] I. Gilboa, D. Schmeidler (1989), “Maxmin Expected Utility with a Non-Unique Prior,” *J. Math Econ*, **18**, 141-153, 1989.
- [11] B. Giroto, S. Holzer (2005), “Representation of Subjective Preferences under Ambiguity,” *J. Math Psychology*, **49**, 372-382, 2005.
- [12] R. Gonzales, G. Wu (1999), “On the Shape of the Probability Weighting Function,” *Cognitive Psychology*, **38**, 129-166.
- [13] L. Hansen, T. Sargent (2001), “Robust Control and Model Uncertainty,” *American Economic Review*, **91**(1), 60-66.
- [14] M. Jahan-Parvar, H. Liu (J-PL, 2012), “Ambiguity Aversion and Asset Prices in Production Economies,” 2012.

- [15] N. Ju, J. Miao (2011), “Ambiguity, Learning, and Asset Returns,” *Econometrica*, forthcoming, 2012.
- [16] P. Klibanoff, M. Marinacci, S. Mukerji (KMM, 2005), “A Smooth Model of Decision Making under Ambiguity,” *Econometrica*, **73**(6), 1849-1892, 2005.
- [17] P. Klibanoff, M. Marinacci, S. Mukerji (2009), “Recursive Smooth Ambiguity Preferences,” *JET*, **144**(3), 930-976, 2009.
- [18] I. Kopylov (2008), “Subjective Probability and Confidence,” UC Irvine Economics Working Paper, 23 May 2008.
- [19] I. Kopylov (2009), “Choice deferral and ambiguity aversion,” *Theoretical Economics* **4**(1), 199-225, 2009
- [20] M. Lindell, J.C. Lu, C.S. Prater (2005), “Household Decision Making and Evacuation in Response to Hurricane Lili,” *Natural Hazards Review*, **6**. 171-179, 2005.
- [21] F. Maccheroni, M. Marinacci, A. Rustichini (2006), “Ambiguity Aversion, Robustness, and the Variational Representation of Preferences,” *Econometrica*, **74**(6), 1447-1498, 2006.
- [22] W. Nordhaus (2007), “The Stern Review on Economics of Climate Change,” *J. Econ. Lit.*, **45**(3), 686-702, 2007.
- [23] W. Nordhaus (2012), “Economic Policy in the Face of Severe Tail Events,” *J. Public Econ. Theory*, **14**(2), 197-219, 2012.
- [24] E.A. Parson (2007), “The Big One: Review of Richard Posner’s ‘Catastrophic Risk and Response,’” *J. Econ. Lit.*, **45**(1), 147-164, 2007.
- [25] J. Quiggin (1982), “A Theory of Anticipated Utility,” *Jnl Econ Behavior & Organization*, **3**, 225-243, 1982.
- [26] D. Schmeidler (1989), “Subjective Probability and Expected Utility without Additivity,” *Econometrica* **57**(3), 571-588, May 1989.
- [27] W.D. Shaw, R.T. Woodward (2008), “Why environmental and resource economists should care about non-expected utility models,” *Resource and Energy Economics*, **30**, 66-89, 2008.
- [28] T. Strzalecki (2011), “Axiomatic Foundations of Multiplier Preferences,” *Econometrica*, **79**(1), 47-74, (2012).
- [29] T. Strzalecki (2012), “Temporal resolution of Uncertainty and Recursive Models of Ambiguity Aversion,” Harvard University working paper, 4 October 2012.
- [30] A. Tversky, D. Kahneman (1992), “Advances in Prospect Theory: Cumulative Representation of Uncertainty,” *Jnl Risk & Uncertainty*, **5**, 297-323, 1992.
- [31] M. Weitzman (2007), “The Stern Review on Economics of Climate Change,” *J. Econ. Lit.*, **45**(3), 703-724, 2007.
- [32] M. Weitzman (2009a), “On Modeling and Interpreting the Economics of Catastrophic Climate Change,” *REStat*, **91**, 1-19, 2009.

- [33] M. Weitzman (2009b), “Risk-Adjusted Gamma Discounting” *NBERw15588*, December 2009.
- [34] J. Werner (2009), “Risk and Risk Aversion When States of Nature Matter,” *Economic Theory* 41(2), 231-246, 2009.