Managing Extinctions*

Massimo Marinacci
Dipartimento di Scienze Economiche
Università di Bologna
Piazza Scaravilli, 2
40126 Bologna, Italy
marinacc@economia.unibo.it

Matthew A.Turner
Department of Economics
University of Toronto
150 St. George Street
Toronto, Ontario M5S 3G7, Canada
mturner@chass.utoronto.ca

August 1998
This version: May 1999

Abstract

A review of the conservation biology literature reveals pervasive uncertainty about the effects of conservation policies on the persistence of populations of endangered species. This uncertainty about population persistence is of two types. Informally, certain features of the world that bear upon population persistence are uncertain, but are drawn from *known* distributions. Other features of the world that bear upon population persistence are uncertain, and are drawn from *unknown* distributions. This paper develops a systematic way to deal with these two types of uncertainty when selecting a conservation policy. In particular, we extend the results of Gilboa and Schmeidler (1989) to the problem of managing the extinction of a single species, and to the more general problem of managing biodiversity.

*We thank Mart Gross, Ian Crane, Ann Zimmerman, and Brian Shuter for their help navigating the conservation biology literature. We also thank Mike Peters, Dale Poirier, Michael Smart, Ralph Winter, Ita Falk, Steven Kohlhagen, and seminar participants at the University of Toronto, Stanford and Harvard for many helpful comments. George Georgopoulis provided valuable research assistance.

1 Introduction

A review of the conservation biology literature reveals pervasive uncertainty about the effects of conservation policies on the persistence of populations of endangered species. This uncertainty about population persistence is of two types. Certain features of the world that bear upon population persistence are uncertain, but are drawn from *known* distributions. Other features of the world that bear upon population persistence are uncertain, and are drawn from *unknown* distributions. This paper proposes that policy makers choose the best policy, conditional on the worst possible realization of those features for which sampling distributions are unknown. That is, we propose a "maximin" decision rule for choosing conservation policies.

There are three arguments in support of our decision rule. First, our rule is based on a description of uncertainty that is more intuitive and more accurate than is possible with standard theory. Second, our rule is consistent with observed choice behavior under very poor information. Third, by extending the results of Gilboa and Schmeidler (1989) to the problem of managing extinctions, we are able to provide an axiomatic foundation for our decision rule. Thus, a policy selected by our decision rule is "rational" in the sense that it is consistent with a simple and plausible set of axioms.

The process of deciding whether a species needs special protection is time consuming, and much of this time is spent resolving scientific uncertainty. Indeed, the Alabama Sturgeon became extinct before scientific uncertainty surrounding its status was resolved (Ando 1997). The decision rule that we propose provides a framework for making decisions on the basis of whatever information is available, however poor. Thus our results could allow policy makers to reduce the large opportunity and fiscal costs of reducing the scientific uncertainty surrounding conservation decisions.

A systematic approach to conservation policy under very poor information may have further benefits. More than half of the federal conservation budget went to only 10 animals, and animals that are longer (head to tail) benefit from more conservation spending (Metrik and Weitzman (1996)). It is plausible that this peculiar allocation of conservation spending is due to the application of *ad hoc* approaches to decision making under uncertainty that differ arbitrarily across species. To the extent that our proposed decision rule systematizes decision making under uncertainty, it may contribute to a better allocation of spending across species.

In addition, the Endangered Species Act mandates particular preferences over extinction densities. We are able to show that, unlike our proposed decision rule, the natural formalization of the mandated preferences violates a "certainty independence axiom". Since this axiom is fundamental to the theory of choice under uncertainty, this raises doubts about the quality of conservation decisions made under this act.

Finally, our decision rule extends to the problem of managing biodiversity as developed in Weitzman (1993), and, Solow, Polasky, and Broadus (1993). Unlike the rules for choosing policies developed by earlier authors, our decision rule explicitly and systematically accounts for the two types of uncertainty typically faced by a policy maker.

1.1 Conservation biology and the "Minimum Viable Population Problem"

To choose one conservation policy over another, we need to know how long a population is likely to persist under the different policies. In the conservation biology literature, this problem is known as the "Minimum Viable Population Problem".[1] The time when a species goes extinct is regarded as stochastic and a solution to Minimum Viable Population Problem involves estimating a distribution of possible extinction times. The most common approach to estimating an extinction density involves the construction of a complex simulation model of the population.[2] These models typically take as inputs: a demographic description of the population, parameters describing death rates and reproductive success for different age cohorts; habitat size; habitat requirements; migration, and the incidence of shocks like drought or fire. Given these inputs, extinction densities are generated numerically by repeated runs of the model. There are of course many variants of this approach, and typically these variants are less structural, e.g., they do not rely on age structured models.[3]

Estimates of densities for extinction times are characterized by the fol-

[1]This nomenclature seems to have arisen with Michael Soule's much cited book, *Viable Populations for Conservation*, Cambridge Univ. Press, Cambridge, 1987.

[2]For examples see: Lindenmayer and Possingham (1996), Middleton and Nisbet (1997), Mills et al (1996), Taylor (1995). We note that in some cases the extinction process is known with a high degree of certainty. For instance, we may be able to infer extinction time very accurately from our predictions about when all habitat will be destroyed. For an example see Harcourt (1995).

[3]For example, Bart (1995) or Middleton and Nisbet (1997).

lowing sorts of uncertainty:

1. Some parameters for simulations may be estimated from field data. These estimates are uncertain, but parameter distributions are known.

2a. Some parameters may be derived from uncertain point estimates taken from one or more field studies (or even made up numbers).

2b. Qualitative features of the models are uncertain. For example: Does the animal migrate? Is population size regulated by food availability or predators?

These types of uncertainty are qualitatively different. While we do not know the true value of parameters that are estimated from field data, we have enough information to construct a probability distribution over these parameters. We follow Knight (1921) and call this type of uncertainty "risk". In contrast, we have considerably less information about parameters derived from point estimates or for qualitative features of the model. Not only we do not know the true state of the world, generally we also do not have enough information to determine a single probability distribution over these states, our "poor" information being compatible with several possible probability distributions. We again follow Knight (1921) and call this type of uncertainty "uncertainty".

1.2 Ellsberg Paradox

The "Ellsberg paradox" (Ellsberg 1961) is a classic and oft-repeated experiment which illustrates the difference between risk and uncertainty and shows that the standard subjective expected utility model[4] is inconsistent with observed preferences over risk and uncertainty.

Suppose there are two urns (I and II), each containing 100 balls, each of which is either white or red. While it is known that in urn I there are 50 white and 50 red balls, there is no information about the number of white and red balls in urn II. In urn I it is natural to assess as equiprobable the drawing of a red ball (R_I) and that of a white one (W_I). For urn II, in contrast, our information is compatible with all probability distributions $P(R_{II})$ and

[4]This standard model of choice under uncertainty was developed by Savage (1954) on the basis of previous works by de Finetti (1931), (1937) and von Neumann and Morgenstern (1944).

$P(W_{II})$ on the drawing of a red (R_{II}) or white (W_{II}) ball. In view of the previous distinction between risk and uncertainty, urn I is characterized by risk and urn II by uncertainty.

A Decision Maker selects a color and one of the two urns, and then draws a ball from the selected urn. If this ball matches a Decision Maker's selected color, the Decision Maker gets \$100. Otherwise the Decision Maker gets nothing.

There are four possible acts: IR = Choose urn I and color R; IW = Choose urn I and color W; IIR = Choose urn II and color R; and IIW = Choose urn II and color W. The state space S consists of four states of nature: s_{RR} = A ball drawn from either urn turns out to be red; s_{RW} = A ball drawn from urn I is red, while a ball from urn II is white; s_{WR} = A ball drawn from urn II is red, while a ball from urn I is white; s_{WW} = A ball drawn from either urn turns out to be white. The correspondence between events and states is: $R_I = \{s_{RR}, s_{RW}\}$, $R_{II} = \{s_{RR}, s_{WR}\}$, $W_I = \{s_{WW}, s_{WR}\}$, $W_{II} = \{s_{WW}, s_{RW}\}$.

A Decision Maker's state contingent payoffs are given in the following table:

	s_{RR}	s_{RW}	s_{WR}	s_{WW}
IR	100	100	0	0
IW	0	0	100	100
IIR	100	0	100	0
IIW	0	100	0	100

Casual empiricism and overwhelming experimental evidence (see, e.g., Camerer (1995)) indicates that most Decision Makers rank these four acts,

$$IR \sim IW \succ IIR \sim IIW. \tag{1}$$

That is, they are indifferent between IR and IW, and between IIR and IIW, but prefer IR and IW to IIR and IIW. Thus choice under risk is different than choice under uncertainty, and in particular, Decision Makers are averse to uncertainty.

Preferences over risk and uncertainty, like those manifested by the ranking (1), cannot be represented with standard subjective expected utility theory. To see this, suppose *per contra* that a Decision Maker ranks actions according to their subjective expected utility. That is, given a probability distribution P over S compatible with a Decision Maker's information, and given a utility function u over the outcomes, a Decision Maker combines them by taking their expectation. For example, the subjective expected utility of the act IR is $u(100)P(R_I) + u(0)P(W_I)$.

If we assume that a Decision Maker has a single probability distribution P over S, then

$$\begin{aligned} IR &\sim IW \Leftrightarrow u(100)P(R_I) + u(0)P(W_I) = u(100)P(W_I) + u(0)P(R_I), \\ IIR &\sim IIW \Leftrightarrow u(0)P(W_{II}) + u(100)P(R_{II}) = u(0)P(R_{II}) + u(100)P(W_{II}), \end{aligned}$$

which imply $P(R_I) = P(W_I)$ and $P(R_{II}) = P(W_{II})$. Hence,

$$P(R_I) = P(W_I) = P(R_{II}) = P(W_{II}) = \frac{1}{2}. \tag{2}$$

On the other hand, by (1), $IR \succ IIR$. Together with (2), this leads to the following contradiction:

$$\begin{aligned} \frac{1}{2}(u(100) + u(0)) &= u(100)P(R_I) + u(0)P(W_I) \\ &> u(0)P(W_{II}) + u(100)P(R_{II}) \\ &= \frac{1}{2}(u(100) + u(0)). \end{aligned}$$

We conclude that the preferences exhibited in the Ellsberg paradox are not compatible with standard subjective expected utility theory. This is not surprising. We observed that the ranking (1) is the result of the relatively poor information about urn II and it seems unlikely that with such limited information a Decision Maker is able to determine a unique probability distribution P over S. On the contrary, it is likely that a Decision Maker considers several probability distributions over S as compatible with his information and that he takes all of them into account when making his decisions.

Now suppose that a Decision Maker faced with poor information about urn II forms a set of prior distributions, all of which are consistent with his information about urn II, and then chooses an urn and a color on the basis of the worst possible prior in this set. This approach allows us to represent the fact that information about urn II is worse than information about urn I, and as we show below, represents observed preferences in the Ellsberg paradox.

Let the set $\mathcal{C}$ of a Decision Maker's priors consist of all P such that

$$P(R_I) = \frac{1}{2}; p_l \leq P(R_{II}) \leq p_h,$$

where $0 \leq p_l < \frac{1}{2} < p_h \leq 1$. That is, the Decision Maker's rich information on urn I delivers a sharp probability $\frac{1}{2}$, while his poor information on urn II makes it possible only to set the lower and upper bounds p_l and p_h. If a Decision Maker then ranks gambles according to the worst prior in the set $\mathcal{C}$ we have:

$$\begin{aligned}
\min_{P \in \mathcal{C}} \left[u(100) P(R_I) + u(0) P(W_I) \right] &= u(100) \frac{1}{2} + u(0) \frac{1}{2}, \\
\min_{P \in \mathcal{C}} \left[u(0) P(R_I) + u(100) P(W_I) \right] &= u(100) \frac{1}{2} + u(0) \frac{1}{2}, \\
\min_{P \in \mathcal{C}} \left[u(100) P(R_{II}) + u(0) P(W_{II}) \right] &= u(100) p_l + u(0) (1 - p_l), \\
\min_{P \in \mathcal{C}} \left[u(0) P(R_{II}) + u(100) P(W_{II}) \right] &= u(100) p_l + u(0) (1 - p_l).
\end{aligned}$$

Since $0 \leq p_l < \frac{1}{2}$, $u(100) \frac{1}{2} + u(0) \frac{1}{2} > u(100) p_l + u(0)(1 - p_l)$. Therefore we have that $IR \sim IW \succ IIR \sim IIW$. This is exactly the ranking (1).

We conclude that a Decision Maker with preferences that can be represented by a maximin rule of the sort demonstrated above is able to systematically account for risk and uncertainty, and, makes choices that are consistent with the preferences over risk and uncertainty that we observe in the Ellsberg paradox.

2 A general description of the extinction problem

Subjective expected utility theory is inconsistent with observed preferences in situations where a Policy Maker faces both risk and uncertainty. Since a choice of conservation policy typically requires a Policy Maker to deal with both risk and uncertainty, standard expected utility theory is not an adequate framework for the choice of conservation policies. This section develops a framework for choosing conservation policies in the presence of both risk and uncertainty.

We follow the conservation biology literature and suppose that extinction times are fundamentally stochastic. This means that the consequence of any choice of conservation policy is a density of extinction times. In light of the uncertainty about estimated extinction densities, we suppose that the relationship between a choice of policy and the resulting extinction density is also random. It follows that the choice of conservation policy is equivalent to the choice of a density over extinction densities. Our problem is to choose between such costly conservation policies.

Let $\mathcal{T} = \{1, 2, 3, ..., t, ...\}$ be the set of all possible extinction times.[5] A probability density f over $\mathcal{T}$ assigns to each subset A of $\mathcal{T}$ a probability. In particular, $f(A) = \alpha \in [0,1]$ reads as follows: "According to f, the probability that extinction occurs at a time $t \in A$ is α". For example, suppose $A = \{2, 4, 6, 8, ...\}$, i.e., A is the set of even numbers, and $\alpha = \frac{1}{2}$. Then $f(A) = \alpha$ means that, according to f, with probability a half extinction will occur at an even time.

We now formally introduce densities. A density f is a set function defined on all subsets of $\mathcal{T}$ such that:

(i) $f(\emptyset) = 0$ and $f(\mathcal{T}) = 1$;

(ii) $0 \leq f(A) \leq 1$ for all $A \subseteq \mathcal{T}$;

(iii) $f\left(\bigcup_{k=1}^{\infty} A_k\right) = \sum_{k=1}^{\infty} f(A_i)$ for all collections $\{A_k\}_{k=1}^{\infty}$ of pairwise disjoint subsets of $\mathcal{T}$.

In particular, (iii) implies that for all $A \subseteq \mathcal{T}$, $f(A) = \sum_{t \in A} f(t)$, and, $\sum_{t=1}^{\infty} f(t) = f(\mathcal{T}) = 1$. Finally, for a given $t' \in \mathcal{T}$, we denote by $f_{t'}$ the density concentrated on t', i.e., $f_{t'}(t) = 0$ if $t \neq t'$ and $f_{t'}(t') = 1$. According to $f_{t'}$ extinction will occur at time t' with probability 1.

For our purposes, a "state of nature" consists of a description of the qualitative features that govern a population's behavior (e.g., whether population density is regulated by predators or food availability), along with a realization of any uncertain parameters in the extinction density (e.g., how many animals there are at $t = 0$, fecundity rates, or the incidence of droughts). That is, given a state of nature, all of a Policy Maker's uncertainty about the extinction density is resolved.

Let $\mathcal{F}$ be the set of all possible densities over extinction times $\mathcal{T}$ and S the set of all relevant states of nature. Subsets of S are called "events". A conservation policy a assigns an extinction density f to each possible state of nature, that is, $a : S \to \mathcal{F}$ is a function from states to extinction densities. $\mathcal{A}$ denotes the set of all conservation policies, that is, the set of all functions $a : S \to \mathcal{F}$. Therefore, because of the fundamentally stochastic nature of

[5] For convenience, we focus on positive integers. At the cost of some slight technical complication, one can consider all non-negative real numbers, i.e., $\mathcal{T} = [0, \infty)$.

extinction times, at each state of nature the consequence of any choice of conservation policy is a density of extinction times.

We close this subsection by defining "constant policies", "simple policies", and "convex combinations". Say that a policy $a \in \mathcal{A}$ is constant if there exists $f \in \mathcal{F}$ such that $a(s) = f$ for all $s \in S$. That is, a induces the same extinction density f for all possible states of nature. For a given $f \in \mathcal{F}$, let a_f denote the constant policy such that $a_f(s) = f$ for all $s \in S$. Moreover, a policy $a \in \mathcal{A}$ is simple if its range $\{a(s) : s \in S\}$ is a finite set. That is, a simple policy induces only finitely many different extinction densities across states.

Given two policies $a, a' \in \mathcal{A}$ and a number $\alpha \in [0,1]$, we denote by $\alpha a + (1-\alpha) a'$ the policy which at state s induces the extinction density $\alpha a(s) + (1-\alpha) a'(s) \in \mathcal{F}$. In other words, $\alpha a + (1-\alpha) a'$ is a policy whose consequences are a mixture of those induced by the two policies a and a'. We say that the policy $\alpha a + (1-\alpha) a'$ is a convex combination of the policies a and a'.

2.1 Preferences

A precursor to a rational choice among costly conservation policies, is the ability to rank conservation policies. Consequently, we require the Policy Maker to rank all conservation policies in a set $\mathcal{A}_{pm} \subseteq \mathcal{A}$, where $\mathcal{A}_{pm}$ is a convex set containing all constant policies.[6]

A Policy Maker's preferences over the policies in $\mathcal{A}_{pm}$ are denoted by $\succeq$, and $a \succeq a'$ reads "the Policy Maker (weakly) prefers a over a'". The preference $\succeq$ is our fundamental notion. Intuitively, a Policy Maker also ranks extinction densities and the extinction times themselves. Such preference orderings on $\mathcal{F}$ and $\mathcal{T}$, denoted respectively $\succeq_{\mathcal{F}}$ and $\succeq_{\mathcal{T}}$ are naturally derived from the basic preference $\succeq$ by using constant policies. The idea is very simple: the constant policy a_f, which always gives f, can be identified with the density f itself, so that one can read $a_f \succeq a_{f'}$ as "the Policy Maker prefers the density f over f'". Given any two densities f and f', we write $f \succeq_{\mathcal{F}} f'$ if $a_f \succeq a_{f'}$. Similarly, if we consider densities $f_{t'}$ concentrated on a single extinction time t', one can read $a_{f_t} \succeq a_{f_{t'}}$ as "the Policy Maker prefers the extinction time t over t'". Given any two extinction times t and t', we

[6] That is, for every $f \in \mathcal{F}$, there exists $a \in \mathcal{A}_{pm}$ such that $a(s) = f$ for all $s \in S$, and if $a, a' \in \mathcal{A}_{pm}$, then $\alpha a + (1-\alpha) a' \in \mathcal{A}_{pm}$ for all $\alpha \in [0,1]$.

write $t \succeq_\tau t'$ if $a_{f_t} \succeq a_{f_{t'}}$.

2.2 Axioms

We assume that a Policy Maker's preferences satisfy the following eight axioms.

A.1 WEAK ORDER: a Policy Maker's preferences are complete and transitive. That is: (i) for all a and a' in $\mathcal{A}_{pm}$, $a \succeq a'$ or $a' \succeq a$; (ii) for all a, a' and a'' in $\mathcal{A}_{pm}$, $a \succeq a'$ and $a' \succeq a''$ imply $a \succeq a''$.

Transitivity is a minimal rationality requirement. Completeness is a more delicate axiom, which excludes the possibility that a Policy Maker cannot rank two policies in $\mathcal{A}_{pm}$. While we expect that this condition is satisfied in many cases of interest, there might well be situations in which the policies are so different that a Policy Maker cannot compare them. Our analysis does not apply to these situations.

A.2 CONTINUITY: For all a, a' and a'' in $\mathcal{A}_{pm}$, if $a \succ a' \succ a''$, then there exist $\alpha, \beta \in (0,1)$ such that

$$\alpha a + (1-\alpha) a'' \succ a' \succ \beta a + (1-\beta) a''.$$

Continuity requires only that no policy be infinitely preferred or despised.

A.3 MONOTONICITY: For all $a, a' \in \mathcal{A}_{pm}$, if $a(s) \succeq_{\mathcal{F}} a'(s)$ for all $s \in S$, then $a \succeq a'$.

Monotonicity[7] requires that if a policy a induces extinction densities $a(s)$ that statewise are preferred to those induced by a', then a Policy Maker prefers a to a'. In other words, if policy a's consequences are state by state better than those induced by a', then a is preferred to a' by a Policy Maker.

A.4 NON-DEGENERACY: There exist $a, a' \in \mathcal{A}_{pm}$ such that $a \succ a'$.

[7] For convenience A.3 is not stated directly in terms of the basic preference $\succeq$, but in terms of $\succeq_m$. The axiom can be formulated in terms of $\succeq$. The same comment applies to axioms A.7 and A.8 below.

This axiom disallows the trivial case where a Policy Maker is indifferent between all elements of $\mathcal{A}_{pm}$.

The first four axioms are common to conventional subjective expected utility and are based on the Anscombe and Aumann's (1963) version of Savage's model. The next two axioms are peculiar to the analysis of Gilboa and Schmeidler (1989).

A.5 UNCERTAINTY AVERSION: For all $a, a' \in \mathcal{A}_{pm}$, if $a \sim a'$ then $\alpha a + (1-\alpha) a' \succeq a$ for all $\alpha \in (0,1)$.

This axiom requires that preferences over acts be quasi-concave, and gives rise to a preference for diversification: given two indifferent policies a and a', a Policy Maker always (weakly) prefers to mix them.

A.6 CERTAINTY INDEPENDENCE: For all $a, a' \in \mathcal{A}_{pm}$, and for all constant acts a'',

$$a \succeq a' if and only if \alpha a + (1-\alpha) a'' \succeq \alpha a' + (1-\alpha) a''$$

for all $\alpha \in (0,1)$.

This weak independence axiom requires that mixing two policies with a common constant policy does not alter the original ordering among the two policies. The common part a'' shared by the policies $\alpha a + (1-\alpha) a''$ and $\alpha a' + (1-\alpha) a''$ is immaterial for their ranking, all that matters is the component in which they differ.[8]

Both A.7 and A.8 reflect the obvious idea that later is better than earlier in extinction times.

A.7 EXTINCTION TIMING: For any two extinction times $t, t' \in \mathcal{T}$, $t \succeq_{\tau} t'$ if and only if $t \geq t'$.

A.8 DOMINANCE: Let $f \in \mathcal{F}$ and $t' \in \mathcal{T}$ be such that $t \geq t'$ for all $t \in supp(f)$. Then $f \succeq_{\mathcal{F}} f_{t'}$.

[8] A.6 is a weak independence axiom because only common parts consisting of constant policies are considered. Standard independence axioms require independence with respect to common general policies (that is, a'' would have been a generic element of $\mathcal{A}_{pm}$, not necessarily constant).

The support $supp\,(f)$ of $f \in \mathcal{F}$ is the set of extinction times to which f assigns a positive probability, that is, $supp\,(f) = \{t \in \mathcal{T} : f\,(t) > 0\}$. According to A.8, if f gives positive probability only to extinction times that occur after t', then the Policy Maker prefers f to a density $f_{t'}$ that gives probability 1 to extinction at time t'.

2.3 Representation of preferences

If the preference $\succeq$ satisfies axioms A.1-A.8, we can provide a simple choice rule that represents preferences over policies. This result is an extension of the representation theorem proved by Gilboa and Schmeidler (1989) to the problem of managing extinctions. In particular, our result generalizes Gilboa and Schmeidler (1989) to allow for an infinite, as opposed to finite set of outcomes, i.e., extinction times, and imposes the additional requirement that "later is better".

We first describe the information available to a Policy Maker. Imagine that a Policy Maker uses his information to formulate beliefs over the relative likelihood of the different events of S. Formally, a belief is a set function P defined on all events of S such that:

1. $P\,(\emptyset) = 0$ and $P\,(S) = 1$;
2. $0 \leq P\,(A) \leq 1$ for all $A \subseteq S$;
3. $P\,(A \cup B) = P\,(A) + P\,(B)$ for all pairwise disjoint sets $A, B \subseteq S$.

In particular, $P\,(A) = \alpha \in [0, 1]$ reads as follows: "According to the Policy Maker, the probability that the true state of nature belongs to A is α". Therefore, P represents the probability that the Policy Maker assigns to the events of S, based on the information available to him. We denote by $\mathcal{P}$ the set of all possible beliefs on S. That is, $\mathcal{P} = \{P \mid P : 2^S \to [0, 1]\}$.

We are now in a position to state the representation theorem. In reading it keep in mind that $a\,(s)$ is an extinction density for each $s \in S$, so that $a\,(s)\,(t)$ is the probability that the density $a\,(s)$ gives to the extinction time t.

Let $\succeq$ be a preference relation on $\mathcal{A}_{pm}$ which satisfies axioms A.1-A.8. Suppose that all the policies in $\mathcal{A}_{pm}$ are bounded, i.e., for all $a \in \mathcal{A}_{pm}$ there exist $f', f'' \in \mathcal{F}$ such that $f' \succeq_{\mathcal{F}} a\,(s) \succeq_{\mathcal{F}} f$ for all $s \in S$, and that $\mathcal{A}_{pm}$ contains all simple policies. Then there exists an increasing utility function

$u : \mathcal{T} \to \mathbb{R}$ and a non-empty convex and compact set $\mathcal{C} \subseteq \mathcal{P}$ of subjective beliefs such that:

(i) for all $t, t' \in \mathcal{T}$, $t \succeq_{\mathcal{T}} t'$ if and only if $u(t) \geq u(t')$;

(ii) for all $f, f' \in \mathcal{F}$, $f \succeq_{\mathcal{F}} f'$ if and only if $\sum_{t=1}^{\infty} u(t) f(t) \geq \sum_{t=1}^{\infty} u(t) f'(t)$;

(iii) for all $a, a' \in \mathcal{A}_{pm}$, $a \succeq a'$ if and only if

$$\min_{P \in \mathcal{C}} \int_S \left[\sum_{t=1}^{\infty} u(t)\, a(s)(t) \right] dP \geq \min_{P \in \mathcal{C}} \int_S \left[\sum_{t=1}^{\infty} u(t)\, a'(s)(t) \right] dP. \quad (3)$$

Moreover, u is unique up to positive linear transformations.

In words, preferences over conservation policies that satisfy A.1-A.8 are equivalent to the following two conditions: (i) there exists a utility function $u : \mathcal{T} \to \mathbb{R}$ which ranks extinction times and such that extinction densities are ranked through their expected utilities; (ii) there exists a set $\mathcal{C}$ of subjective beliefs such that each policy a is ranked by first considering the subjective expected utility $\int_S [\sum_{t=1}^{\infty} u(t)\, a(s)(t)]\, dP$ associated with each belief, and then focusing on the worst one.

Theorem 2.3 is the fundamental result of the paper. It provides a framework for managing extinctions which systematically accounts for different qualities of information. A non-singleton set of priors, $\mathcal{C}$, and the minimization rule (3) reflect both a Policy Maker's limited degree of confidence in his beliefs (several beliefs being compatible with his information) and his dislike for this type of uncertainty. This innate attitude towards uncertainty is formalized by axiom A.5 and we believe that most individuals have this negative attitude toward vagueness of beliefs.

Technically, Theorem 2.3 extends the original result of Gilboa and Schmeidler (1989) to densities having infinite supports (their analysis considers only finite support densities). However, the main contribution of Theorem 2.3 is in showing how to apply their setting to conservation issues.

2.4 A two state example

We now illustrate Theorem 2.3 with a simple two state example, where $S = \{s_1, s_2\}$. In this case, each policy can be directly identified with the pair

(f_1, f_2) of induced extinction densities. If Theorem 2.3's representation holds, by (3) the policy's payoff can be summarized by the pair

$$\left(\sum_{t=1}^{\infty} u(t) f_1(t), \sum_{t=1}^{\infty} u(t) f_2(t) \right) \in \mathbb{R}_+^2.$$

For convenience, we will denote such pairs simply by $(x_1, x_2) \in \mathbb{R}_+^2$. Constant policies are characterized by $x_1 = x_2$.

Theorem 2.3's preferences over policies are then given by

$$V(x_1, x_2) = \min_{p \in [\underline{p}, \overline{p}]} p x_1 + (1-p) x_2, \tag{4}$$

where p is the probability of s_1, and $0 \leq \underline{p} \leq \overline{p} \leq 1$. Solving the minimization problem (4), we have:

$$V(x_1, x_2) = \begin{cases} \overline{p} x_1 + (1 - \overline{p}) x_2 & \text{if } x_1 < x_2, \\ \underline{p} x_1 + (1 - \underline{p}) x_2 & \text{if } x_1 \geq x_2. \end{cases}$$

If the set $[\underline{p}, \overline{p}]$ is a singleton, the indifference map is:

Figure 1.

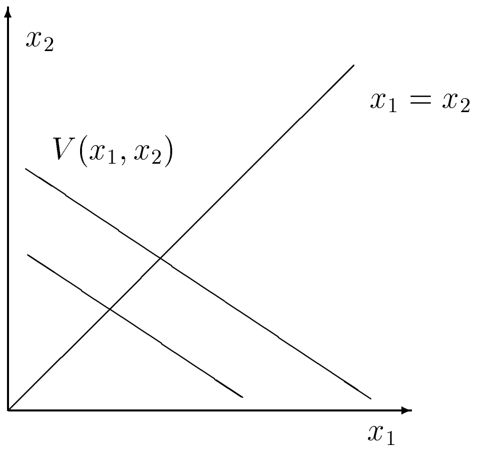

If the set $[\underline{p}, \overline{p}]$ is not a singleton, then the indifference curves kink at the certainty line $x_1 = x_2$:

Figure 2.

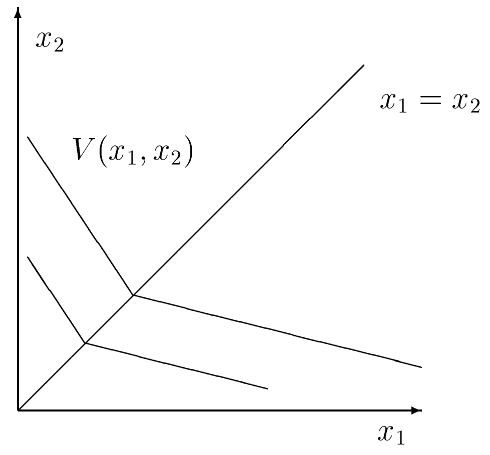

As the set $[\underline{p}, \overline{p}]$ approaches $[0,1]$, the kink's angle approaches ninety degrees. Proposition 2.4 will show formally that singleton sets of priors occur if and only if A.5 holds with indifference, that is, if $a \sim a'$ implies $\alpha a + (1-\alpha) a' \sim a$ for all $\alpha \in (0,1)$. Therefore, kinks can only arise when A.5 holds strictly, that is, when a true preference for diversification occurs.

Kinks reflect uncertainty aversion. As the set of priors becomes larger and, consequently, as indifference curves become more sharply kinked, a Policy Maker is more willing to exchange increases in expected payoff, for payoffs which are closer to the certainty line and hence more nearly the same in all states. In this way, a Policy Maker exhibits uncertainty aversion, an attitude which is also reflected by the fact that, as the sets of priors become larger and larger, a Policy Maker requires lower and lower certainty equivalents for each given policy.[9]

We now explain why axioms A.1-A.6 require an indifference map of the sort illustrated in the figures, and hence minimization over a set of priors. A.1-A.4 are standard axioms playing a standard role. For example, A.1 guarantees that indifference curves do not intersect, while A.3 ensures that a Policy Maker is better off on higher indifference curves.

A.5 is quasi-concavity axiom and, therefore, implies convex indifference curves. As to A.6, consider two indifferent policies, a and a'. A.6 requires

[9]For a given policy a, its certainty equivalent c is a constant policy such that $a \sim c$.

that, for any constant policy c,

$$\alpha a + (1 - \alpha) c \sim \alpha a' + (1 - \alpha) c.$$

This is true in particular if $a' = c$, so that

$$\alpha a + (1 - \alpha) c \sim c.$$

Thus, A.6 requires that all indifference curves take a very particular form: If a non-constant policy and a constant one lie on the same indifference curve, so do all their convex combinations. In sum:

$$\begin{aligned} A.5 &\implies Convex indifference curves, \\ A.6 &\implies Piecewise linear indifference curves. \end{aligned}$$

Together, A.5 and A.6 then require that all indifference maps take the form illustrated in Figures 1 or 2, and so the axioms require an indifference map of the sort produced by minimizing over a set of priors.

We now formalize the intuition that without preference for diversification we should get back to a standard subjective expected utility model (i.e., indifference maps like in Figure 1).

Under the hypotheses of Theorem 2.3, A.5 holds with indifference if and only if the set $\mathcal{C}$ is a singleton. That is, there is a single probability P such that, for all $a, a' \in \mathcal{A}_{pm}$, $a \succeq a'$ if and only if

$$\int_S \left[\sum_{t=1}^{\infty} u(t) a(s)(t) \right] dP \geq \int_S \left[\sum_{t=1}^{\infty} u(t) a'(s)(t) \right] dP.$$

Intuitively, Policy Makers who are uncertainty neutral behave in the same way regardless of the confidence they have in their beliefs. Such a Policy Maker would form a unique prior for urn II in our version of the Ellsberg paradox.

Finally, we note that "uncertainty loving" Policy Makers are characterized by the following dual version of A.5: For all $a, a' \in \mathcal{A}_{pm}$, if $a \sim a'$ then $\alpha a + (1 - \alpha) a' \preceq a$ for all $\alpha \in (0, 1)$.

One can show that, if in Theorem 2.3 we replace A.5 with this dual version, then $a \succeq a'$ if and only if

$$\max_{P \in \mathcal{C}} \int_S \left[\sum_{t=1}^{\infty} u(t) a(s)(t) \right] dP \geq \max_{P \in \mathcal{C}} \int_S \left[\sum_{t=1}^{\infty} u(t) a'(s)(t) \right] dP,$$

for all $a, a' \in \mathcal{A}_{pm}$. That is, a Policy Maker selects the best policy given the best possible resolution of his uncertainty.

3 The Optimization Problem and Some Special Cases

Based on Theorem 2.3, we can write a Policy Maker's choice of conservation policy as an optimization problem,

$$\max_{a\in\mathcal{A}_{pm}}\left[\min_{P\in C}\int_S\left(\sum_{t=1}^{\infty}u(t)\,a(s)(t)\right)dP-c(a)\right], \tag{5}$$

where $c(a)$ is the cost associated with action $a\in\mathcal{A}_{pm}$.[10]

This statement of a Policy Maker's problem assumes that a Policy Maker's preferences are quasi-linear in species and income. While this is a restrictive assumption, it is appropriate for the case we have in mind: A Policy Maker with a small conservation budget managing one of many species. If we are interested in a conservation budget large enough to impact the marginal rate of substitution between species and money, then a different formulation of the problem is appropriate.

We now consider different specifications of the Policy Maker's problem, and describe the resulting optimization problems.

Example 1. Unique prior: Conventional subjective expected utility is a special case of the present setting in which $\mathcal{C}$ is a singleton $\{P\}$.[11] In this case, $a\succeq a'$ if and only if

$$\int_S\left[\sum_{t=1}^{\infty}u(t)\,a(s)(t)\right]dP\geq\int_S\left[\sum_{t=1}^{\infty}u(t)\,a'(s)(t)\right]dP,$$

so that problem (5) becomes:

$$\max_{a\in\mathcal{A}_{pm}}\left[\int_S\left(\sum_{t=1}^{\infty}u(t)\,a(s)(t)\right)dP-c(a)\right].$$

[10] The set of feasible policies available to the Policy Maker will generally be much smaller than $\mathcal{A}_{pm}$, and will probably not contain all the constant and simple policies. To accommodate this, we may assign a very large cost to infeasible policies, or else specify the feasible subset of $\mathcal{A}_{pm}$. To avoid introducing further notation, we use the former strategy.

[11] We note that Goodman (1998) provides a more extensive treatment of this special case.

This version of the Policy Maker's problem arises when a Policy Maker perceives no uncertainty in his estimate of extinction densities. As noted above, this may reflect either high quality information about the extinction density, or a Policy Maker's uncertainty neutrality.

Example 2. Maximal uncertainty: This version of the Policy Maker's problem describes a situation in which a Policy Maker perceives the quality of the information to be maximally poor, and provides a simple example where the use of the maximin rule is normatively compelling. Let $\mathcal{C}$ consist of all possible beliefs over S (that is, $\mathcal{C} = \mathcal{P}$). Since $\mathcal{C}$ contains all priors which assign probability one to a single state of nature, it is easy to show that

$$\min_{P \in \mathcal{C}} \int_S \left[\sum_{t=1}^{\infty} u(t)\, a(s)(t) \right] dP = \inf_{s \in S} \left[\sum_{t=1}^{\infty} u(t)\, a(s)(t) \right],$$

so that problem (5) becomes:

$$\max_{a \in \mathcal{A}_{pm}} \left[\inf_{s \in S} \left(\sum_{t=1}^{\infty} u(t)\, a(s)(t) \right) - c(a) \right]. \tag{6}$$

Example 3. Favorite prior: We now present a case intermediate between the above two polar cases and which is convenient analytically. For given "favorite prior" $P_0 \in \mathcal{P}$ and $\alpha \in [0,1]$, set $\mathcal{C} = \{\alpha P_0 + (1-\alpha) P : P \in \mathcal{P}\}$. Intuitively, P_0 is a belief that, given his information, stands out in a Policy Maker's view. However, because his information is not good enough, he only gives to P_0 a weight α, leaving the remaining $(1-\alpha)$ unassigned. Clearly, the better a Policy Maker's information, or the less his aversion to uncertainty, the higher α will be. If $\alpha = 1$, we get back to conventional subjective expected utility, while with $\alpha = 0$ we end up with $\mathcal{C} = \mathcal{P}$, the case of maximally poor information.

For this specification it is easy to show that

$$\min_{P \in \mathcal{C}} \int_S \left[\sum_{t=1}^{\infty} u(t)\, a(s)(t) \right] dP = \alpha \int_S \left[\sum_{t=1}^{\infty} u(t)\, a(s)(t) \right] dP_0 + (1-\alpha) \inf_{s \in S} \left[\sum_{t=1}^{\infty} u(t)\, a(s)(t) \right],$$

and so problem (5) becomes:

$$\max_{a \in \mathcal{A}_{pm}} \left[\alpha \int_S \left(\sum_{t=1}^{\infty} u(t)\, a(s)(t) \right) dP_0 + (1-\alpha) \inf_{s \in S} \left(\sum_{t=1}^{\infty} u(t)\, a(s)(t) \right) - c(a) \right],$$

which is a combination of the maximin rule and subjective expected utility with respect to P_0.

Example 4. Threshold preferences: One implication of Theorem 2.3 is that axioms A.1-A.8 are logically consistent with each other. Therefore, if a Policy Maker likes the axioms, it makes sense to try to make decisions consistent with the axioms. Conversely, a Policy Maker should try to avoid making decision inconsistent with the axioms. This turns out to be a quite powerful tool for checking the "rationality" of policy decisions. To illustrate this, we consider a formalization of the preferences over extinction densities that are mandated by the Endangered Species Act.

The Endangered Species Act requires that species be partitioned into those that are "endangered" or "threatened", and hence receive legal protection, and those that are not endangered or threatened and do not receive legal protection. In the conservation biology literature there is considerable interest in "threshold" type preferences over extinction densities as a natural way to systematize this mandated process of partitioning the set of species into those that are endangered or threatened and those that are not.[12]

For the sake of illustration, say that a species is endangered if its extinction densities result in a greater than 5% probability of extinction in 100 years. These threshold preferences $\succeq_{\mathcal{F}}$ over $\mathcal{F}$ have the following numerical representation $U : \mathcal{F} \to \mathbb{R}$,

$$U(f) = \begin{cases} -k & \text{if} \quad \sum_{t=1}^{100} f(t) > .05 \\ 0 & \text{else} \end{cases} .$$

That is, a Policy Maker pays a welfare penalty k whenever there is a 5% or greater probability of extinction before 100 years have passed.

On first glance, threshold preferences seem like an intuitive and natural way to formalize the process of deciding whether or not species are endangered. These preferences, however, do not satisfy the conditions of Theorem 2.3. In particular, they violate axiom A.6.

To see that A.6 is violated, consider the extinction densities f, f', $f'' \in \mathcal{F}$ such that:

$$\begin{aligned} f(10) &= .01, f(110) = .99; \\ f'(10) &= .04, f'(110) = .96; \\ f''(10) &= 1. \end{aligned}$$

[12]See Taylor (1995) for a discussion of these preferences and further references.

We have $U(f) = U(f')$, i.e., $f \sim_m f'$. Let $\alpha = \frac{98}{100}$. According to A.6, we should then have $\alpha f + (1-\alpha) f'' \sim_m \alpha f' + (1-\alpha) f''$. But, $\sum_{t=1}^{100} (\alpha f(t) + (1-\alpha) f''(t)) < 0.05$ and $\sum_{t=1}^{100} (\alpha f'(t) + (1-\alpha) f''(t)) > 0.05$, so that $\alpha f + (1-\alpha) f'' \succ_m \alpha f' + (1-\alpha) f''$, thus violating A.6. This means that threshold preferences over extinction densities do not satisfy the independence axiom, and, consequently, do not have an expected utility representation.

In sum, threshold-type preferences violate the independence axiom. Since this is a basic tenet of rational choice under uncertainty, in our opinion, this casts serious doubt on the appropriateness of using threshold-type preferences as a guide to conservation policy.

There is another way in which Theorem 2.3 can help Policy Makers. If a Policy Maker knew the types of information that typically cause people to perceive uncertainty (e.g. Urn II of our Ellsberg paradox), and also knew that people are uncertainty averse, then by making more conservative decisions when faced with information that causes uncertainty (and vice verse) a Policy Maker can make people in the economy better-off . This suggests the need for research to characterize the types of information that lead people to perceive uncertainty.

Example 5. Stationary preferences and Scientific priors: Up until now, we have taken our primitive to be a Policy Maker's preferences over policies. While this is convenient analytically, for practical problems we are more likely to start with preferences over extinction times, $\succeq_{\mathcal{T}}$, and densities, $\succeq_{\mathcal{F}}$. Furthermore, in applications a Policy Maker often has a scientific description of the uncertainty surrounding extinction densities. For convenience, we call this information a "scientific prior". Though priors over states of nature (the set $\mathcal{C}$ in Theorem 2.3) are necessarily subjective, when a Policy Maker has full confidence in the scientific description of uncertainty, he may equate "scientific priors" with his subjective priors. For example, suppose that a Policy Maker is concerned about tomorrow's whether and he receives a forecast saying that "the probability of rain tomorrow is between 20% and 30%". If a Policy Maker has full confidence in such a scientific description of weather uncertainty, his subjective prior will coincide with the forecast, that is, "scientific priors" and subjective priors will coincide.

In this case, Theorem 2.3 gives us a rule for making decisions consistent with scientific estimates of the uncertainty surrounding estimated extinction densities: A Policy Maker should choose the best policy, conditional on the worst resolution of scientific uncertainty.

For the sake of illustration, let k denote the cost of an extinction, whenever it occurs, and suppose that we are willing to discount the future at some constant rate. In this case, the utility associated with a particular extinction time is $u(t) = -k\delta^t$, and we can represent numerically our preferences over extinction densities as follows:

$$f \succeq_{\mathcal{F}} f' if and only if -\sum_{t=0}^{\infty} k\delta^t f(t) \geq -\sum_{t=0}^{\infty} k\delta^t f'(t). \tag{7}$$

That is, a Policy Maker prefers extinction densities that result in lower expected discounted extinction costs.

Suppose that a Policy Maker is uncertain about only two aspects of the extinction density, the precise number of breeding pairs in the current population, and whether the population size is regulated by predators or food availability. Suppose that a Policy Maker has a unique prior over the number of breeding pairs, but is maximally uncertain about the mechanism by which population density is regulated. More formally, let N be the set of possible values for the number of breeding pairs, and $\mathcal{C}^N = \{P^N\}$ a Policy Maker's (singleton) set of priors over this set. Let $D = \{D^p, D^f\}$ denote the set of two possible mechanisms which regulate population density, and $\mathcal{C}^D = \{(P^{D^p}, 1 - P^{D^p}) : 0 \leq P^{D^p} \leq 1\}$ a Policy Maker's set of priors over this set. Thus, a Policy Maker's problem involves both risk and uncertainty.

The set of states of nature S is the Cartesian product of $\{D^p, D^f\}$ and N, so that any particular state of nature $s \in S$ is a pair (D^i, n), with $n \in N$ and $i \in \{p, f\}$. A Policy Maker's set of priors over these states is given by the Cartesian product of the two sets of priors, so that $\mathcal{C} = \mathcal{C}^D \times \mathcal{C}^N$. More formally, $\mathcal{C} = \mathcal{C}^D \times \mathcal{C}^N$ consists of all $P : 2^S \to [0,1]$ such that $P(D^i, n) = P^{D^i} P^N(n)$ for all $(D^i, n) \in S$.

A Policy Maker must choose between conservation policies a that, in each state of nature, induce Poisson extinction densities over extinction times,

$$a(s)(t) = \frac{\lambda_a(s)^t e^{-\lambda_a(s)}}{t!},$$

where λ_a is the function from states of nature into the positive real numbers that is associated with conservation policy a. Let $\mathcal{A}_{pm}$ be the smallest convex set containing all such "Poisson" policies available to a Policy Maker, and all constant and simple policies.

Let $\succeq$ be a preference over $\mathcal{A}_{pm}$ satisfying A.1-A.8 and which induces the preferences over densities given in equation (7). Then, by Theorem 2.3, a

Policy Maker's preferences over the set of "Poisson" conservation policies are represented by the following function:

$$V(a)=\min_{P\in\mathcal{C}}\left[\sum_{(D^i,n)\in S}\left(\sum_{t=0}^{\infty}k\delta^t\frac{\lambda_a(D^i,n)^t e^{-\lambda_a(D^i,n)}}{t!}\right)P\left(D^i,n\right)\right] \tag{8}$$

Since $P\left(D^i,n\right)=P^{D^i}P^N\left(n\right)$, using some simple algebra we can simplify the right hand side of (8) to

$$\begin{aligned}&\min_{P^D\in\mathcal{C}^D}\sum_N\left(P^{D^p}\sum_{t=0}^{\infty}k\delta^t\frac{\lambda_a(D^p,n)^t e^{-\lambda_a(D^p,n)}}{t!}\right.\\&\left.+\left(1-P^{D^p}\right)\sum_{t=0}^{\infty}k\delta^t\frac{\lambda_a(D^f,n)^t e^{-\lambda_a(D^f,n)}}{t!}\right)P^N\left(n\right).\end{aligned}$$

Given these preferences it is trivial to write down a Policy Maker's optimization problem corresponding to (5).

This discussion provides a simple example of the way that different qualities of information about important determinants of extinction densities can be incorporated systematically into management decisions. Generalizing to more complicated and realistic problems occurs in the natural way: If there is uncertainty over more than two parameters, then the set of possible priors expands; if the family of extinction densities under consideration is constructed by simulation, then the Policy Maker's problem is solved numerically.

4 Managing biodiversity

Our analysis has so far considered the way that poor quality information about extinction densities can be used in management decisions for individual species. A closely related problem is the management of sets of species, that is, the problem of managing biodiversity.[13] In this section we extend our analysis to this more general problem.

[13] The problem of managing biodiversity has been addressed in the literature, e.g., Weitzman (1992), (1993), and Solow, Polasky and Broadus (1993). However, the existing research is primarily interested in developing ways to measure the diversity of given sets of species, and not with the uncertainty surrounding extinction densities.

Suppose we are interested in a set $I = \{1, ..., I\}$ of species, where for convenience I denotes both the set and its last element. Let $\mathcal{T}^I = \mathcal{T} \times \cdot \cdot \cdot \times \mathcal{T}$ be the set of extinction times for the I species, with generic element $(t_1, ..., t_I) \in \mathcal{T}^I$. A joint extinction density f over $\mathcal{T}^I$ assigns to each vector $(t_1, ..., t_I) \in \mathcal{T}^I$ a probability $\alpha \in [0, 1]$. For example, $f(t_1, ..., t_I) = \frac{1}{2}$ reads "with probability one half extinction will occur at time t_1 for species 1, at time t_2 for species 2,..., at time t_I for species I".[14]

Joint extinction densities also satisfy properties (i)-(iii) of page 8, where A and A_k are subsets of $\mathcal{T}^I$ instead of $\mathcal{T}$. Let $\mathcal{F}^I$ be the set of all possible joint extinction densities. A conservation policy a is now a function from states to joint extinction densities, that is, from S into $\mathcal{F}^I$. In particular, $a(s)(t_1, ..., t_I)$ is a joint extinction density, i.e., it is the probability that at state $s \in S$ the policy a gives to extinction occurring at $(t_1, ..., t_I)$. Note that if the set I is a singleton the problem of managing biodiversity reduces to the problem considered earlier.

Up to this change in the interpretation of policies and densities, almost all notions introduced in Section 2 are meaningful in the more general "biodiversity" context. In particular, this is the case for the induced preferences $\succeq_{\mathcal{T}}$ and $\succeq_{\mathcal{F}}$, and for axioms A.1-A.6. The only two exceptions are the statements of axioms A.7 and A.8, which critically depend on the single dimensionality of $\mathcal{T}$. However, these axioms can be easily adapted to the more general setting. In particular, A.7 now takes the form:

A.7* For any given $i \in I$ and any two extinction times $t', t'' \in \mathcal{T}$,

$$(t_1, ..., t_{i-1}, t', t_{i+1}, ..., t_I) \succeq_{\mathcal{T}} (t_1, ..., t_{i-1}, t'', t_{i+1}, ..., t_I)$$

if $t' \geq t''$.

In words, A.7* says that, *ceteris paribus*, a Policy Maker prefers a vector of extinction times that gives a later extinction time to a given species $i \in I$. That is, if all other species' extinction times are unchanged, a later extinction time is better for species i.

As to A.8, we now have:

A.8* Let $f \in \mathcal{F}^I$ and $(t'_1, ..., t'_I) \in \mathcal{T}$ be such that $(t_1, ..., t_I) \succeq (t'_1, ..., t'_I)$ for all $(t_1, ..., t_I) \in supp(f)$ then $f \succeq_{\mathcal{F}} f_{(t'_1, ..., t'_I)}$. Conversely, if $(t_1, ..., t_I) \preceq (t'_1, ..., t'_I)$ for all $(t_1, ..., t_I) \in supp(f)$, then $f \preceq_{\mathcal{F}} f_{(t'_1, ..., t'_I)}$.[15]

[14] With a slight abuse of notation, we write $f(t_1, ..., t_I)$ in place of $f((t_1, ..., t_I))$.

[15] Here, $supp(f) = \{(t_1, ..., t_I) \in \mathcal{T}^I : f(t_1, ..., t_I) > 0\}$.

This dominance axiom says that a joint density which gives positive probability only to vectors of extinction times preferred ("dispreferred", resp.) to $(t'_1, ..., t'_I)$, must then be preferred ("dispreferred", resp.) to the degenerate density $f_{(t'_1,...,t'_I)}$ that gives probability 1 to $(t'_1, ..., t'_I)$.

4.1 Representation of preferences

Theorem 2.3 can be readily adapted to the more general setting of biodiversity. In particular, we have:

Let $\succeq$ be a preference relation on a set $\mathcal{A}_{pm}$ of "biodiversity" policies $a : S \to \mathcal{F}^I$ which satisfies axioms A.1-A.6, and A.8*. Suppose that all the policies in $\mathcal{A}_{pm}$ are bounded, and that $\mathcal{A}_{pm}$ contains all simple policies. Then there exists a utility function $u : \mathcal{T}^I \to \mathbb{R}$, and a non-empty convex and compact set $\mathcal{C} \subseteq \mathcal{P}$ of subjective beliefs such that:

(i) for all $(t_1..., t_I), (t'_1..., t'_I) \in \mathcal{T}^I$, $(t_1..., t_I) \succeq_{\mathcal{T}} (t'_1..., t'_I)$ if and only if $u(t_1..., t_I) \geq u(t'_1..., t'_I)$;

(ii) for all $f, f' \in \mathcal{F}^I$, $f \succeq_{\mathcal{F}} f'$ if and only if $\sum_{(t_1...,t_I)\in\mathcal{T}^I} u(t_1..., t_I) f(t_1..., t_I) \geq \sum_{(t_1...,t_I)\in\mathcal{T}^I} u(t_1..., t_I) f'(t_1..., t_I)$;

(iii) for all $a, a' \in \mathcal{A}_{pm}$, $a \succeq a'$ if and only if

$$\begin{aligned} & \min_{P\in\mathcal{C}} \int_S \left[\sum_{(t_1...,t_I)\in\mathcal{T}^I} u(t_1..., t_I)\, a(s)(t_1..., t_I) \right] dP \\ \geq\ & \min_{P\in\mathcal{C}} \int_S \left[\sum_{(t_1...,t_I)\in\mathcal{T}^I} u(t_1..., t_I)\, a'(s)(t_1..., t_I) \right] dP. \end{aligned}$$

Moreover, u is unique up to positive linear transformations.

(iv) u is increasing in each argument if and only if A.7* holds.[16]

[16] Remember that $u(\cdot)$ is increasing in each argument if

$$u(t_1, ..., t_{i-1}, t', t_{i+1}, ..., t_I) \geq u(t_1, ..., t_{i-1}, t'', t_{i+1}, ..., t_I)$$

whenever $t' \geq t''$

The interpretation of Theorem 4.1 is similar, *mutatis mutandis,* to that of Theorem 2.3. In particular, examples 1, 2, and 3 are still meaningful for Theorem 4.1.

We can also relate Theorem 4.1 to recent work that seeks to characterize and apply diversity indices to the problem of selecting a conservation policy (e.g., Weitzman 1992 and 1993). To do this, we first define the "survival function". Given $(t_1, ..., t_I) \in \mathcal{T}^I$, the survival function $Q_{t_1,...,t_I} : \mathcal{T} \to I$ is defined as follows: $Q_{t_1,...,t_I}(t) = \{i \in I : t_i > t\}$.[17] A diversity function $D : 2^I \to \mathbb{R}$ is a real-valued set function which assigns a value to each set of species according to how "valuable" the Policy Maker thinks such a set is. We assume that $D(\emptyset) = 0$.

Following Weitzman (1993), suppose that a Policy Maker's preferences over extinction vectors are given by $u : \mathcal{T}^I \to \mathbb{R}$, where u is the present value of the diversity for any given vector of extinction times. That is,

$$u(t_1, ..., t_I) = \sum_{t=1}^{\infty} D\left(Q_{t_1,...,t_I}(t)\right)\delta^t \tag{9}$$

for all $(t_1, ..., t_I) \in \mathcal{T}^I$.

Now, let $\succeq$ be a preference over a set $\mathcal{A}_{pm}$ of "biodiversity" policies $a : S \to \mathcal{F}^I$. If $\succeq$ satisfies A.1-A.6 and A.8*, and if it induces the preferences over densities given in equation (9), then, by Theorem 4.1, a Policy Maker's ranking of conservation policies to preserve biodiversity is represented by the following criterion:

$$\min_{P \in \mathcal{C}} \int_S \left(\sum_{(t_1 ..., t_I) \in \mathcal{T}^I} \left[\sum_{t=1}^{\infty} D\left(Q_{t_1,...,t_I}(t)\right)\delta^t \right] a(s)(t_1 ..., t_I) \right) dP. \tag{10}$$

If we assign a cost $c(a)$ to each conservation policy available to a Policy Maker, the problem of choosing a conservation policy to preserve biodiversity reduces to the following maximization problem:

$$\max_{a \in A} \left[\min_{P \in \mathcal{C}} \int_S \left(\sum_{(t_1 ..., t_I) \in \mathcal{T}^I} \left[\sum_{t=1}^{\infty} D\left(Q_{t_1,...,t_I}(t)\right)\delta^t \right] a(s)(t_1 ..., t_I) \right) dP - c(a) \right].$$

[17] In other words, $Q_{t_1,...,t_I}$ lists which species in I still survive at time t, given a vector of extinction times $(t_1, ..., t_I)$. For example, if $I = \{1, 2\}$ and $(t_1, t_2) = (60, 40)$, then $Q_{60,40}(30) = \{1, 2\}$, $Q_{60,40}(50) = \{1\}$, $Q_{60,40}(100) = \emptyset$.

Weitzman (1992), (1993) and Solow, Polasky and Broadus (1993) show how diversity functions can be derived from basic dissimilarity distances among species, and then provide competing rules for selecting conservation policies. Our model, and the resulting decision rule (10), extends their analysis by explicitly and systematically accounting for the quality of a Policy Maker's information about the joint extinction density.

4.2 Monotonicity

Say that a diversity index D is monotone if for any $A, B \subseteq I$, we have $D(A) \leq D(B)$ whenever $A \subseteq B$. That is, a biodiversity index is "monotone" if adding species to a set cannot decrease the value of the index. Interestingly, if a Policy Maker's preferences over vectors of extinction times are given by Theorem 9, A.7* and monotone diversity functions are formally equivalent. The intuition behind this result follows. If A.7 holds, the present value of diversity must increase if any species survives another day. This can only occur if the diversity index always increases with the number of species. Conversely, if the diversity index is monotone, whenever a species survives for another day, discounted diversity must also increase.

The utility function in (9) satisfies A.7* if and only if D is monotone.

5 Conclusion

We have presented a decision rule which allows a policy maker to make systematic comparisons of different conservation policies given both risk and uncertainty. Our rule has three important features. First, a Policy Maker following our rule behaves in a way that is consistent with observed choice behavior under risk and uncertainty. He chooses the policy that is best, conditional on the worst possible resolution of uncertainty. Second, our analysis shows that such conservative behavior is justified by a simple and plausible axiomatic description of preferences over conservation policies, thus making transparent the behavioral foundation of our rule. Finally, our rule allows a Policy Maker to make decisions that effectively take into account the quality of the information available to him, in particular his degree of confidence in such an information. Indeed, as Frank Knight remarked in his 1921 classic work, "the action which follows upon an opinion depends as much upon the amount of confidence in that opinion as it does upon the favorableness of the opinion itself." We believe that these three features of our decision rule make it quite attractive as a basis for the choice of conservation policies.

Our analysis also has immediate policy implications for the Endangered Species Act. In particular, the Endangered Species Act mandates particular preferences over extinction densities. We are able to show that the natural formalization of the mandated preferences does not satisfy a basic tenet of rational choice under uncertainty.

Finally, we are able to extend our framework to the problem of managing biodiversity, that is, to the problem of managing the extinction of a set of species. For biodiversity as well, we generate a decision rule that explicitly and systematically accounts for the two types of uncertainty typically faced by a Policy Maker.

We close by observing that a limitation of our analysis is that it supposes that the value of an extinction be known with certainty. In fact, there is great uncertainty about the value of a species. Intuitively, one could think of dealing with this problem by, for example, enlarging the state space or by allowing several possible values and then taking their lower envelope. A rigorous approach to the problem, however, appears difficult and remains a subject for further research.

6 Proofs

Proof of Theorem 2.3. Because of A.7, axiom A.8 can be written as follows: let $f \in \mathcal{F}$ and $t' \in \mathcal{T}$ such that $t \succeq_{\mathcal{T}} t'$ for all $t \in supp(f)$, then $f \succeq_{\mathcal{F}} f_{t'}$. Now, suppose $f \in \mathcal{F}$ and $t' \in \mathcal{T}$ are such that $t \preceq_{\mathcal{T}} t'$ for all $t \in supp(f)$. We want to show that $f \preceq_{\mathcal{F}} f_{t'}$ in this case. By A.7, $t \preceq_{\mathcal{T}} t'$ for all $t \in supp(f)$ implies that $t \leq t'$ for all $t \in supp(f)$. Hence, $supp(f)$ is a finite set. Let $f_1, f_2, f_3, f_4 \in \mathcal{F}$ be any extinction densities such that $f_1 \succeq_{\mathcal{F}} f_2$ and $f_3 \succeq_{\mathcal{F}} f_4$. A simple application of A.6 shows that for every $\alpha \in (0,1)$ it holds that $\alpha f_1 + (1-\alpha) f_3 \succeq_{\mathcal{F}} \alpha f_2 + (1-\alpha) f_4$. By applying a finite number of times this fact, we get $\sum_{t \in supp(f)} f(t) f_t \preceq_{\mathcal{F}} f_{t'}$ if $t \preceq_{\mathcal{T}} t'$ for all $t \in supp(f)$. But $f(t) = \sum_{t \in supp(f)} f(t) f_t(t)$ for all $t \in \mathcal{T}$, and so $f \preceq_{\mathcal{F}} f_{t'}$, as wanted. We conclude that, for all $t \in supp(f)$,

$$t \succeq_{\mathcal{T}} t' \Longrightarrow f \succeq_{\mathcal{F}} f_{t'}$$

$$t \preceq_{\mathcal{T}} t' \Longrightarrow f \preceq_{\mathcal{F}} f_{t'}.$$

This implies the dominance axiom A.4* of Fishburn (1982). On the other hand, we have that: (i) $\mathcal{F}$ contains all densities over the power set $2^{\mathcal{T}}$, and so it satisfies the structural condition A0.1 of Fishburn (1982), (ii) $\succeq_{\mathcal{F}}$ is a weak order by A.1, it is continuous by A.2, and it satisfy the independence axiom by A.6. Therefore, all hypotheses of Theorem 3 p. 28 of Fishburn (1982) are satisfied, and so there exists a bounded real-valued function $u : \mathcal{F} \to \mathbb{R}$ such that

$$f \succeq_{\mathcal{F}} f' if and only if \sum_{t=1}^{\infty} u(t) f(t) \geq \sum_{t=1}^{\infty} u(t) f'(t) \tag{11}$$

for all $f, f' \in \mathcal{F}$. Moreover, u is unique up to positive linear transformations (clearly, u is increasing by A.7).

Define $U(f) = \sum_{t=1}^{\infty} u(t) f(t)$. By (11), the function U represents the ordering $\succeq_{\mathcal{F}}$. Moreover, it is easy to see that U is unique up to positive linear transformations and that it is an affine function on $\mathcal{F}$, i.e., $U(\alpha f + (1-\alpha) f') = \alpha U(f) + (1-\alpha) U(f')$ for all $f, f' \in \mathcal{F}$.

Since $\succeq_{\mathcal{F}}$ is essentially the restriction of $\succeq$ on $\mathcal{F}$, we can summarize in the following claim what we have proved till now:

Claim. There exists an affine $U : \mathcal{F} \to \mathbb{R}$ such that for all $f, f' \in \mathcal{F}$: $f \succeq f'$ iff $U(f) \geq U(g)$. Furthermore, U is unique up to a positive linear transformation.

Lemma 3.1 p. 145 of Gilboa and Schmeidler (1989), the first step of their derivation, shows the existence of a similar function on the subspace of $\mathcal{F}$ consisting of all densities having a finite support.[18] Our claim extends their lemma to the entire $\mathcal{F}$, i.e., to the space of all densities, with finite or infinite support. Our extended version can then replace their original Lemma 3.1 and be used as a first step of a more general derivation dealing with all densities, with finite or infinite supports. In fact, it can be checked that all subsequent steps of Gilboa and Schmeidler's (1989) derivation (i.e., Lemmas 3.2, 3.3, 3.4, and 3.5) hold also under our Claim. It suffices to keep in mind that in our setting Σ is the power set $2^{\mathcal{T}}$, $L(\geqq)$ is the set of all bounded policies (there are no measurability issues since we are dealing with a power set), and that, using their notation on p. 150, we have $L_0 \subseteq \mathcal{A}_{pm} \subseteq L(\geqq)$ (and so their Remark on p. 150 applies to $\mathcal{A}_{pm}$).

[18]In their notation, u is our U, and Y is the subspace of $\mathcal{F}$ of all finite support densities, i.e., $f \in Y$ if $f(t) \neq 0$ for at most a finite number of points of time $t \geq 1$.

If $\mathcal{F}$ consisted of finite support densities, Theorem 2.3 would reduce to the original Gilboa and Schmeidler's (1989) representation result. Since we just showed how to modify their derivation to take care of the presence in $\mathcal{F}$ of densities with infinite supports, the proof is complete. □

Proof of Proposition 2.4. Together with A.6, A.5* implies the following independence condition: For all $a, a', a'' \in \mathcal{A}_{pm}$,

$$a \succeq a' if and only if \alpha a + (1-\alpha) a'' \succeq \alpha a' + (1-\alpha) a'', \tag{12}$$

for all $\alpha \in (0,1)$. In fact, notice that, as easily checked, a'' has a certainty equivalent c. Hence, by A.5*,

$$\begin{aligned} \alpha a + (1-\alpha) a'' &\sim \alpha a + (1-\alpha) c, \\ \alpha a' + (1-\alpha) a'' &\sim \alpha a' + (1-\alpha) c. \end{aligned}$$

By A.6, we conclude that $\alpha a + (1-\alpha) a'' \succeq \alpha a' + (1-\alpha) a''$, as wanted. Axioms A.1-A.4, and the independence condition (12), are those used in the standard Anscombe-Aumann model, and so the probability P representing the likelihood relation is unique. □

Proof of Theorem 4.1. The dominance axiom A.8* is the exact counterpart in the context of biodiversity of the dominance axiom A.4 in Fishburn (1982). Therefore, we can invoke again his result to obtain an affine function on $\mathcal{F}^I$ such that $f \succeq_{\mathcal{F}} f'$ if and only if

$$\sum_{(t_1...,t_I) \in \mathcal{T}^I} u(t_1...,t_I) f(t_1...,t_I) \geq \sum_{(t_1...,t_I) \in \mathcal{T}^I} u(t_1...,t_I) f'(t_1...,t_I)$$

for all $f, f' \in \mathcal{F}^I$. By proceeding as in the proof of Theorem 2.3, we complete the proof. □

Proof of Proposition 4.2. Suppose D is monotone. Let

$$(t_1, ..., t_{i-1}, t', t_{i+1}, ..., t_I), (t_1, ..., t_{i-1}, t'', t_{i+1}, ..., t_I) \in \mathcal{T}^I$$

with $t' \geq t''$. Then $Q_{t_1,...,t_{i-1},t'',t_{i+1},...,t_I}(t) \subseteq Q_{t_1,...,t_{i-1},t',t_{i+1},...,t_I}(t)$ for all $t \geq 1$, and so, by monotonicity of D,

$$D\left(Q_{t_1,...,t_{i-1},t'',t_{i+1},...,t_I}(t)\right) \subseteq D\left(Q_{t_1,...,t_{i-1},t',t_{i+1},...,t_I}(t)\right)$$

for all $t \geq 1$, so that

$$\sum_{t=1}^{\infty} D\left(Q_{t_1,...,t_{i-1},t',t_{i+1},...,t_I}\left(t\right)\right)\delta^t \geq \sum_{t=1}^{\infty} D\left(Q_{t_1,...,t_{i-1},t'',t_{i+1},...,t_I}\left(t\right)\right)\delta^t,$$

as wanted.

As to the converse, suppose $u\left(t_1,...,t_I\right)$ is increasing in each argument. For any $A \subseteq I$, let $\left(t_1^A,...,t_I^A\right) \in \mathcal{T}^I$ be defined as follows: $t_i^A = 1$ if $i \notin A$ and $t_i^A = 2$ if $i \in A$. Then, $\sum_{t=1}^{\infty} D\left(Q_{t_1^A,...,t_I^A}\left(t\right)\right)\delta^t = \delta D\left(I\right) + \delta^2 D\left(A\right)$. Suppose $A, B \subseteq I$ are such that $A \subseteq B$. Then, each component of $\left(t_1^A,...,t_I^A\right)$ is less or equal than the corresponding component of $\left(t_1^B,...,t_I^B\right)$. Hence, applying a finite number of times the fact that $u\left(t_1,...,t_I\right)$ is increasing in each argument, we get

$$\delta D\left(I\right) + \delta^2 D\left(A\right) = u\left(t_1^A,...,t_I^A\right) \leq u\left(t_1^B,...,t_I^B\right) = \delta D\left(I\right) + \delta^2 D\left(B\right),$$

and so $D\left(A\right) \leq D\left(B\right)$, which shows that D is monotone. $\square$

References

[1] **A. W. Ando**. "Delay on the path the endangered species list: Do costs and benefits matter?," Resources for the Future, DP 97-43 (1997).

[2] **F.J. Anscombe and R.J. Aumann**. "A definition of subjective probability," *Ann. Math. Stat.* 34, 199-205, (1963).

[3] **J. Bart.** "Amount of suitable habitat and viability of Northern Spotted Owls," *Conservation Biology* 9, 943-946, (1995).

[4] **G. E. Belovsky.** "Extinction models and mammalian persistence", in *Viable Populations for Conservation* (M. E. Soule, Ed.), Cambridge Univ. Press, Cambridge, 1987.

[5] **C. Camerer**. "Individual Decision Making," in *Handbook of Experimental Economics* (J.H. Kagel, and A.E. Roth, Eds.), Princeton Univ. Press, Princeton, 1995.

[6] **B. de Finetti**. "Sul Significato Soggettivo della Probabilità," *Fund. Math.* 17, 298-329 (1931). (English translation in: B. de Finetti, *Probabilità ed Induzione*, CLUEB, Bologna, 1993).

[7] **B. de Finetti.** "La Prevision: Ses Lois Logiques, Ses Sources Subjectives," *Ann. Inst. Poincaré* 7, 1-68, (1937). (English translation in: H. E. Kyburg and H. E. Smokler, *Studies in Subjective Probability*, Wiley, New York, 1964).

[8] **P. Fishburn**. *The foundations of expected utility*, Kluwer, Dordrecht, 1982.

[9] **I. Gilboa and D. Schmeidler.** "Maxmin expected utility with non-unique prior," *Journal of Mathematical Economics* 18, 143-153 (1989).

[10] **D. Goodman.** "Population Viability Analysis: The Risk Assessment Paradigm," Biology Department, Montana State University, Mimeo (May 1998).

[11] **A.H. Harcourt.** "Population viability estimates: Theory and practice for a wild Gorilla population," *Conservation Biology* 9, 134-142, (1995).

[12] **F.H. Knight**. *Risks, Uncertainty, and Profit*, Houghton Mifflin, Boston, 1921.

[13] **D. B. Lindenmayer and H. P. Possingham.** "Ranking conservation and timber management options for Leadbeater's Possum in Southeastern Australia using Population Viability Analysis," *Conservation Biology* 10, 235-251, (1996).

[14] **A. Metrik and M. L. Weitzman.** "Patterns of Behavior in Endangered Species Preservation," *Land Economics 72(1), 1-16, (*1996).

[15] **D. A. J. Middleton and R. M. Nisbet.** "Population persistence time: Estimates, models, and mechanisms," *Ecological Applications* 7, 107-117, (1997).

[16] **L. Scott Mills, S. G. Hayes, C. Baldwin, M. J. Wisdom, J. Citta, D. J. Mattson, and K. Murphy.** "Factors leading to different viability predictions for a Grizzly Bear data set," *Conservation Biology* 10, 863-873, (1996).

[17] **L. J. Savage.** *The foundations of statistics*, Wiley, New York, 1954.

[18] **A. Solow, S. Polasky, and J. Broadus.** "On the Measurement of Biological Diversity," *Journal of Environmental Economics and Management*, 24(1), 60-68 (1993).

[19] **B. L. Taylor.** "The reliability of using population viability analysis for risk classification of species," *Conservation Biology* 9, 235-251, (1995).

[20] **J. von Neuman and O. Morgenstern**. *Theory of games and economic behavior*, Princeton Univ. Press, 1944.

[21] **M.L. Weitzman.** "On diversity," *Quarterly Journal of Economics* 107, 363-406 (1992).

[22] **M.L. Weitzman.** "What to preserve? An application of diversity theory to crane conservation," *Quarterly Journal of Economics* 108(1), 157-183 (1993).