On The Welfare Aspects of Species Conservation Policy: A Stochastic Dominance Approach

Greg Hunter
Cal Poly Pomona

Steve Polasky
University of Minnesota

Draft: July 29, 2007

Protecting valuable ecosystem services from the unintended consequences of human activity is a concern shared by many. The robust interaction of a diverse multitude of species inhabiting the ecosystem ensures the abundant provision of these services. Everyone concerned about the threat of human activity to biodiversity should be equally concerned that resources devoted to species conservation are wisely used. In this paper, we analyze the economic efficiency aspects of the species conservation policies in a cost-effectiveness framework.

Characterizations of economic efficiency typically yield insights for improving policy. Our analysis provides a method for comparing the efficiency implications of different conservation planning approaches in the literature and allows us to offer policy guidance. A technique for ranking planning approaches according to their impact on a well-defined measure of social welfare is absent from the literature. The absence of a structure for comparing the welfare aspects of various approaches to policy points to a larger gap in the literature in public goods provision under uncertainty. Our aim is to extend the work of others in the field of welfare theory toward providing a structure that is more amenable to application.

We expose a model of collective choice under uncertainty in the presence of a public good. The model was inspired by our concern for articulating the welfare aspects of species conservation policy, but thorough consideration of the issues uncovered necessitated an approach that has significance well beyond the domain of conservation planning. We begin by defining a stochastic process that generates levels of a jointly consumed good, the provision of which concerns a benevolent planner and a collection of rational individuals, I. Next, we assume a non-paternalistic role for the planner who identifies social efficiency by appealing to the Pareto criterion. To place our work in a context that relates directly with the existing literature on conservation planning, we limit the scope of our analysis to focus on *cost-effective* plans, which are *cost-constrained* Pareto efficient. We find that identifying an efficient allocation at this level of generality requires appealing to the mechanism design literature, which is an impractical guide for policy. Subsequently, we seek additional restrictions that enable useful characterizations of the efficient set of conservation plans. These characterizations limit the degree to which individuals may disagree about the desirability of different policy outcomes and the probability of the outcomes occurring. Eventually, when disagreement about the ranking of outcomes and probability beliefs are completely eliminated, a representative agent framework arises and characterizations of the efficient set become very sharp and fragile (in a sense to be explained later).

Measuring sticks for the rational comparison of jointly consumed risky prospects must account for the ways in which personal values may differ under uncertainty. Because these ways are many, the economic theory of rational individual choice under uncertainty, *or expected utility theory*, is marked by a level of generality that leads to little practical advice. *A fortiori,* practical approaches to collective choice under uncertainty rarely have transparent connections to economically meaningful concepts like *Pareto Optimality, the Samuelson Criterion, or the Compensation Criterion*. It could easily be argued that the generality required for an adequate description of choice under uncertainty precludes the existence of practical measuring sticks for policy analysis. In what follows, we articulate sufficient conditions on the choice set and individual preferences allowing us to implement an approach to collective choice under uncertainty that is both practical and well grounded in economic theory.

Species conservation planning is our leading example for illustrating the usefulness of this method. We have chosen to couch our discussion in terms of conservation planning for two reasons. First, the study was prompted by our desire to understand the economic consequences of conservation policy with an eye toward identifying the relative merits of various proposals. Second, it turns out that conservation planning conforms naturally to the restrictions required by our approach.

However in exposing our innovation, we have found the clarity of our exposition calls for us to confront many of the more general issues of welfare comparisons when dealing with public goods provision under uncertainty. Thus, our technique has policy implications for a broad class of problems that adhere to a few of the features that make conservation planning an especially good example. In particular, we require that there be a well-defined set of mutually exclusive and exhaustive social states, the realization of which can be determined with relative easy and consensus. We also require a common-knowledge, objective probability measure over the social states be induced by each policy.

Our approach borrows its structure from the literature dealing with rational choice under uncertainty. As such, we regard extinction risk as a "game against nature" or *game of chance*. In the absence of human conservation efforts, each species, $n \in N \subset \mathbb{N}$, has a *null policy* survival probability represented by a probability space, $\phi_0 \equiv (\Omega_0, \vartheta_0, P_0)$. Each element of the sample space, Ω_0, is an *N*-dimensional vector, or *collection*, indicating the survival status of each species. ϑ_0 is the associated σ-algebra and P_0 is the associated probability measure, which gives the survival probability for each collection. We assume a planner with preferences that are *collectively rational* in the sense of Pareto. We also assume that the planner has an individualistic and non-paternalistic view of social welfare. To begin identifying characteristics of policies preferred by such a planner necessitates consideration of individual preferences over the set of feasible policies, $\phi \in \Phi \subset \mathbb{N}$. Each policy, $\phi \equiv (\Omega_\phi, \vartheta_\phi, P_\phi)$, is a probability space defined by a probability measure, P_ϕ, over the various collections of species that are potential outcomes of the extinction process. We assume the probability measure for each policy is unique for convenience. Note that our analysis would be easier if we could also assume that all feasible policies share a common sample space. In general, however, policy outcomes, Ω_ϕ , will include both species collections and a distribution of aggregate wealth, which may affect the ranking of survival outcomes. To avoid this complication, we limit the scope of our analysis slightly by focusing on identifying characteristics of *cost-constrained Pareto efficient* policies holding the distribution of wealth constant. Hereafter, we refer to such policies as "cost-effective," which implies the cost of the conservation plan, K_ϕ, has been determined exogenously, individual payments are sunk, and that the policy satisfies a well-defined sense of collective rationality. Hence, from this point forward policies will share a common sample space, $\Omega = \Omega_\phi, \forall \phi \in \Phi$.

In this setting, each feasible conservation plan, $\phi \equiv (\Omega, \vartheta, P_\phi)$, is a game of chance like roulette, dice, or coin tossing. In the economics of uncertainty literature, there is no consensus on how to define rational preferences over games of chance. Rather, it is customary to take a "consequetialist" view that rational agents derive utility not from the game itself, but from the consequences of the game (e.g., prizes from winning). Here, we follow custom and define a space of consequences for the set of cost-constrained feasible policies, C^i, for each individual, $i \in I \subset \mathbb{N}$. Following the mainstream economics of uncertainty literature, we assume that rational choice under uncertainty is consistent with the axioms of expected utility. Thus, analytically an individual's choice over a set of conservation policies is much like trying to decide whether to play dice or roulette.[1]

Expected utility analysis of games of chance usually focuses on a closely related issue. Typically analysts have confined their attention to examining the choice of agents confronted with selecting an optimal strategy for playing a particular game of chance. Think of choosing whether to place a bet on "red" or "black" in a game of roulette. There is a corresponding choice made by conservation planners that we ignore for the moment so that we can focus on choosing an optimal plan, which is more aptly characterized as a choice between games. Still, this is the proper place in our exposition for laying the groundwork for the subsequent consideration of the issues raised by this decision 'sub-problem.' The uncertainty literature commonly refers to the set of strategies available to an individual playing a particular game as "acts." We define, A_ϕ as the space of acts for policy ϕ. Each act, $a \in A_\phi$, is a mapping from ϕ into C^i (i.e., a random variable). Each act induces a probability measure for each individual, $\Pi^i_{a|\phi}$, on the events defined by C^i. At this point in our analysis, we assume that for each individual, i, there is a single optimal act, a^{i*}, that is independent of the individual's optimal policy, ϕ^{i*}, which allows to use the unambiguous notation, Π^i_ϕ, when discussing individual preferences over policy alternatives. Making use of the games of chance metaphor, we can think of a policy choice as a choice between similar games (i.e., with a common set of available actions). For example, we might think of someone trying to decide whether to play roulette in Monte Carlo or Las Vegas. Assuming the wheels may be biased, this is a decision of sincere concern. In our case, we have to insist that the differences between the roulette wheels are not so significant that the player would change betting strategies.

General Necessary and Sufficient Conditions for Cost-Effectiveness

In this setting, expected utility theory is an appropriate vehicle for discussing an individual's preferences over policy, which is the first step to an economically meaningful analysis. Appealing to the expected utility theorem, we can assume the existence of a function, $U^i(C^i_\phi) = \int v^i(c) d\Pi^i_\phi$, that expresses the individual's preferences over policy consequences.

---

[1] This characterization is made to aid in the flow of the exposition. Strictly, speaking we could enlarge our space of consequences and treat the "game" choice as a strategy choice, but doing so prevents us from making a valuable point.

Notice, $U(\bullet)$ is linear in the probabilities of the various consequences and that $v^i(c)$ is defined uniquely to linear affine transformations for each individual. Recognizing the policy consequences of concern are multi-dimensional in that for each individual they influence both the collection of surviving species, ω, and the amount of a state-contingent numeraire, y, we can define the individual's willingness-to-pay, or *reservation price,* r_ϕ^i, for a policy, ϕ, that is different from the null policy as the r that satisfies $\int v^i(\omega, y-r)d\Pi_\phi^i - \int v^i(\omega, y)d\Pi_0^i = 0$.[2] In this setting, the Samuelson condition for optimal public goods provision implies that optimal policies will be members of the set: $\Phi^* \equiv \{\phi^* : \sum_i r_\phi^i \leq \sum_i r_{\phi^*}^i, \forall \phi \in \Phi\}$. In other words, an optimal policy generates a higher summed willingness-to-pay than any other feasible policy alternative.

> **Defintion 1:** Under uncertainty, a feasible cost-constrained policy for providing a public good, $\phi \equiv (\Omega, \vartheta, P_\phi)$, is *cost-effective* if and only if $\phi \in \Phi^* \equiv \{\phi^* : \sum_i r_\phi^i \leq \sum_i r_{\phi^*}^i, \forall \phi \in \Phi\}$, where Φ represents the set feasible cost-constrained policies.

As articulated, the *Samuelson Condition* is both necessary and sufficient for optimal public goods provision under uncertainty. However, it suffers from a defect in that no operational form of the Condition exists. Green & Laffont (1980) and others have extensively studied the problem of truthfully eliciting an individual's preferences for public goods provision, but magnitude of the government payments required to incentivize truth telling are prohibitive in most cases.

Other notions of cost-effectiveness

We can derive an equivalent expression for Φ^* by considering the Bergson-Samuelson social welfare function (SWF) $L = \sum_I \lambda^i \int v^i(\omega, y)d\Pi_\phi$. In this case, we can express the Pareto efficient set of policies as:

$$\Phi \supset \Phi^* \equiv \{\phi^* : \forall i \in I, \exists \lambda^i \geq 0 \text{ and strictly for at least one, } \phi^* = \arg\max L(\phi; \lambda)\}.$$

Since the integral, L, is linear in the vector of welfare weights, λ, this expression provides the foundation for a social ranking of policy *outcomes*, as well as a social ranking of policies. Our first definition of cost-effectiveness (i.e., Definition 1) obscures the importance of the social ranking of the policy outcomes, which is essential to an economically meaningful definition of cost-effectiveness. Put another way, a particular policy may be "cost-effective"

---

[2] There is slight abuse of notation by introducing ω directly into the agent's preference scaling function. Since, ω is a meaningful description of the social state there should not be any confusion. In the case of a conservation plan, ω would include the survival status of each species.

in achieving one set of policy priorities and completely “ineffective” at achieving another. This realization is embodied in the Bergson-Samuelson SWF and provides the basis for narrowing the scope of our analysis of cost-effectiveness.

> **Defintion 2:** Under uncertainty, a feasible cost-constrained policy for providing a public good, $\phi \equiv (\Omega, \vartheta, P_\phi)$, is *conditionally cost-effective* if and only if $\phi \in \Phi_\lambda^* \equiv \{\phi_\lambda^* : \phi_\lambda^* = \arg\max L(\phi, \lambda),$ where is a λ vector of specific welfare weights$\}$
> .

This definition says that a *conditionally* cost-effective policy will maximize expected social welfare for a particular SWF. However, agreeing on the relative social value of the various policy outcomes leads to little, or no, advice for policymakers because it requires specific and detailed information about the preferences of every individual affected by the various proposals, as well, as a weighting of those individuals' influence on the policy action. Note that any policy chosen for a particular set of welfare weights has no guarantee of being chosen (i.e., maximizing expected social welfare) if there is even an arbitrarily small change in the welfare weights that preserves the underlying social preference relation (i.e., ordinal ranking of social states).

Given the fragility of conditional cost-effectiveness to small changes in the SWF, one might reasonably wonder whether it would be worth the effort to gather the data to rank the various policy outcomes. Our approach assumes that an ordinal ranking of the social states being considered will be relatively easy for policymakers to generate.

If policymakers do the work of ranking the social states (i.e., policy outcomes) there is an economically meaningful definition of cost-effectiveness that can be employed to identify policies robust to small changes in SWF weights that preserve the social preference ordering. In what follows, we use the following definition of cost-effectiveness and show that it achieves the goal of being well connected to welfare theory and at the same time yields practical guidance for policy. The definition makes use of an idea for ranking random variables called *first-order stochastic dominance (FSD).* Next, we summarize the idea of FSD and some of the key results used in the remainder of our paper.

FSD has been used to robustly rank random variables with *monetary* payoffs for sometime (Rothschild and Stiglitz (1970)). To our knowledge, this is the first time the concept of FSD as been used in a welfare context. To extend the concept to a welfare setting, we have to assume a *linear order*[3] on the social states exists, which is the essential feature of monetary payoffs of random variables that makes them a good candidate for FSD analysis. For example, the property of a linear order in the context of monetary payoffs means that \$2 is always better than \$1 regardless of the decision makers underlying preferences. It is also true that a linear order does not allow for a “tie” in the ranking between distinct social states. When a linear order exists on the state space, robust comparisons of random variables are accomplished by examining the differences between their cumulative distribution functions

[3] Formally, a linear order is complete, transitive and antisymetric.

(CDF). When a random variable X stochastically dominates another random variable Y, the CDF of variable X, F_X lies everywhere below (weakly) that of Y, F_Y. When X *FSD* Y, we will use the notation $X >> Y$. Rothschild and Stiglitz (1970) show that when $X >> Y$, every decision maker with monotone preferences (i.e., that agrees with the linear ordering) will prefer X to Y. In the policy context, this implies higher social welfare. Note, that Rothschild and Stiglitz (1970) also show that $X >> Y \Rightarrow E(X) > E(Y)$, but the converse does not hold. This means that merely choosing policy to maximize expected social welfare using a particular weighting of the social states does not ensure selection of a stochastically dominant policy. So, FSD is a stricter requirement than maximizing expected social welfare.

When a social ranking of the policy outcomes is chosen, we know that if a policy, $\hat{\phi}$, stochastically dominates an alternative, then $\hat{\phi}$ will generate higher expected SWF for all social welfare functions that preserve the underlying social preference relation.

Defintion 3: Under uncertainty, a feasible cost-constrained policy for providing a public good, $\phi \equiv (\Omega, \vartheta, P_\phi)$, is *locally cost-effective* if and only if $\hat{\phi} \in \hat{\Phi} \equiv \{\phi : \forall i \in I, \exists \lambda^i \geq 0$ and strictly for at least one, $\hat{\phi} = \arg\max L(\phi; n(\lambda, \varepsilon))\}$, where $n(\lambda, \varepsilon)$ is a neighborhood of welfare weights that preserves the ordering of the social states implicit in λ.

The following relationship between the definitions of cost effectiveness:

Theorem 1: $\Phi^* \subset \Phi^*_\Gamma$, where $\Phi^*_\Gamma = \{\phi^*_\lambda : \phi^*_\lambda \in \cup_\lambda \Phi^*_\lambda\}$

Proof:

Suppose this is not true. Then, there exists a policy that maximizes the sum of the willingness to pay, but does not maximize a weighted sum agents' expected utilities, where the weights may be arbitrarily chosen. This implies that there is at least one individual whose willingness to pay cannot be increased by choosing a different feasible alternative, but the SWF is not maximized for any feasible weighting of expected utility functions. Suppose we choose a SWF that applies zero weight to all other agents and a strictly positive weight to the agent in question. If the agent's expected utility is increasing in the numeraire, then her willingness to pay will increase if an alternative with higher expected utility is selected!

Theorem 1 says the set of policies that satisfies the Samuelson condition in Definition 1 is a subset of the set of policies that maximize a Bergson-Samuelson welfare function. This is an important fact that we alluded to earlier because the Bergson-Samuelson welfare weights that correspond to those set of policies that satisfy the Samuelson condition induce a linear order on the set of policy outcomes. This ordering is of special significance because it contains information about social preference relation implicit in the Samuelson condition.

Theorem 2: $\hat{\Phi} \subset \Phi^*_\Gamma$, where $\Phi^*_\Gamma = \{\phi^*_\lambda : \phi^*_\lambda \in \cup_\lambda \Phi^*_\lambda\}$

Proof: This is trivial when the set of FSD policies is empty and true by the definition of $\hat{\Phi}$ otherwise.

Theorem 3: If λ^* is a set of welfare weights that corresponds to the maximization of a SWF, $L(\phi, \lambda^*)$, whose $\arg\max$,ϕ^*, satisfies the Samuelson Condition: $\sum_i r^i_\phi \leq \sum_i r^i_{\phi^*}, \forall \phi \in \Phi$, r^i_ϕ satisfies $\int v^i(\omega, y-r)d\Pi^i_\phi - \int v^i(\omega, y)d\Pi^i_0 = 0$,

then $\hat{\Phi}^* \subset \Phi^*$, where $\hat{\Phi}^* \equiv \{\hat{\phi} : \hat{\phi} = \arg\max L(\phi; n(\lambda^*, \varepsilon))\}$ and $n(\lambda^*, \varepsilon)$ is a neighborhood of welfare weights that preserves the ordering of the social states implicit in λ^*.

Proof: This is trivial when the set of policies $\hat{\Phi}^*$ is empty. Otherwise, recall that these policies are FSD for the ranking of policy outcomes implicit in λ^*. Since FSD implies higher expected social welfare and willingness to pay is increasing in expected utility, our conclusion is obtained.

Corollary: If no FSD policy is a member of Φ^*, then the remaining non-trivial elements must *undominated* in the first-order stochastic sense.

Theorem 3 and its corollary articulate our approach to narrowing the search for, and in some cases identifying, cost-effective policies under uncertainty. There are three steps. First, determine a ranking of the social states of concern that is consistent with the underlying social preference relation. This ranking should be a linear order on the set of social states, though it does not have to assign a cardinal value to each outcome of the policy. The states themselves should correspond to a *partition* of the underlying sample space. This ensures that the states are mutually exclusive and exhaustive, which fulfills a technical requirement for a well-defined probability measure. Second, use the probability measure on the set of social states for each policy to construct a cumulative distribution function (based on some cardinal ordering that is consistent with the ranking of the social states).[4] Lastly, identify any candidate policy with a first-order stochastically dominant CDF. Assuming that policies generate a unique CDF, there will be at most one policy whose CDF dominates all other alternatives. If no such policy exists, then any policy with a CDF that is dominated in the first-order sense by the CDF of some other feasible policy should be eliminated. The remaining set of policies, those with undominated CDFs, are conditionally cost-effective for a set of welfare weights that are consistent with the implied ordering of social states in step 1.

It should be clear that this procedure is general and could be used for welfare comparison of any set of policies that generate a probability distribution over a set of mutually exclusive

[4] This is another minor technical detail to ensure that we are comparing the distribution functions of random variables; however, our analysis could extend to the comparison of *pseudo-random* variables that merely rely on the social ranking described.

and exhaustive social states. The key to our approach is that it relies on the ability of policymakers to generate an ordering of the social state with approximate knowledge of the social preference relation.

A species focused approach to conservation planning

Most approaches to conservation planning formulate objectives aimed at increasing the survival probabilities of particular species. We refer to such approaches as being "species centered." Our method of welfare comparison is well suited to evaluating these approaches. The feature of species centered approaches to conservation planning that makes them well suited to our methodology is that the species survival outcomes form a mutually exclusive and exhaustive sample space over which to define the probability measures induced by the policy alternatives being considered. Additionally, when it comes to ranking the set of policy outcomes, there are a small number of considerations and they are easy to specify.

A Simple Example

Choosing a reserve site is intended to increase the survival probabilities of species who take refuge in its bounds. In large reserve sites a relatively small number of species have direct aesthetic, cultural, or immediate commercial value. There is, however, a general concern for species preservation motivated by a desire to conserve valuable sources of as-yet-undiscovered information.

Yet despite careful planning, reserve sites are an inherently imperfect preservation strategy, which generate a vector of survival probabilities for affected species that typically are not reasonably close to unity. The uncertainty can be because of imperfect species location data, which are then used in the reserve selection process or because the habitat preservation strategy is insufficient to protect a particular species from all causes of extinction (e.g., disease and predation).

So, the exercise of rational reserve site selection amounts to choosing among lotteries for information and ecosystem services, where "ecosystem services" is interpreted to include all use and non-use values associated with biodiversity. Accordingly, the optimal reserve site corresponds to the most valuable lottery. There is a difficulty in comparing the value of species lotteries (i.e. reserve sites) because the payoff (i.e. outcome of the lottery) includes information and preferences for information tend to be non-quasi-concave, though not always. Still, any robust ranking of reserve sites should be invariant to whether the planner's preferences for information are risk-loving, risk-neutral, or risk-averse. The only ranking method for lotteries that meets the invariance criterion is first-order stochastic dominance (Rothschild and Stiglitz, (1970)).

In the following simple example, we demonstrate how to use first-order stochastic dominance to rank reserve sites based on their informative potential. The method appropriately ranks conservation programs for any preference relation that is monotonic in species. Weitzman's (1998) method for measuring information content in a set of species is used to generate a Blackwell (1953) index for each set. The Blackwell index has the convenient property that it is increasing in the value of information whenever preferences are monotone in consumption. In the example, it will be assumed that the conservation planner has a fixed budget. There are three potential reserve sites {L, J, K} each of which contains two affected species. The planner can preserve any two of the three sites. Without loss of generality, assume that species in the unprotected site go extinct. Also assume that each site contains only two distinct species identified by a single genetic difference.

We think of the reserve sites as sets whose cardinality measures information content: L = {x,z}; J = {x,y}; K = {x,y}, where x, y, and z are used to identify the affected species. For concreteness assume that if a site is chosen as a reserve, then affected species in the site have a survival probability of $p = 0.25$ and that extinction events in each reserve are probabilistically independent. There are three possible conservation plans: $L \cup K; L \cup J; K \cup J$. All of these plans have the possibility of saving all species affected by the particular reserve designation and the possibility of having all affected species go extinct. In the case that all species are saved, the reserve's informative potential is at an upperbound. Designate the *value* of the maximal information content for each plan as $\bar{\phi}_{LUK}; \bar{\phi}_{LUJ}$; and $\bar{\phi}_{KUJ}$, respectively. These measure of the value of information in the sense of Blackwell (1953), which assumes that each gene is an experiment that may yield useful information about future states of the world. A higher Blackwell ranking implies that every monotonic agent will receive greater expected utility than she would if exposed to a lower ranked information set. Note that $\bar{\phi}_{LUK} = \bar{\phi}_{LUJ} > \bar{\phi}_{KUJ}$. However, this ranking of the conservation plans is based only on the best outcome.

A thorough consideration of all the possible outcomes and their implied probabilities allows for a robust ranking of the plans based on the criterion of first-order stochastic dominance. To begin, the cumulative distribution function (CDF) for the value of information must be constructed. Because every gene in the example represents a single experiment whose probability of success is independent of the other experiments conducted, the Blackwell ranking (i.e., value of information) is increasing in the number of preserved genes (cardinality of the preserved set). A value of information that increases in the number of preserved genes facilitates comparisons of the CDFs generated by different conservation plans. The following tables list the events and associated probabilities for each feasible plan in the example:

Event Probabilities for Plan $K \cup J$

x-Events	y-Events	Joint Event
Pr(x) = 0.5625	Pr(y) = 0.5625	Pr($K \cup J$ = {x,y}) = 0.3164
Pr(~x) = 0.4375	Pr(y) = 0.5625	Pr($K \cup J$ = {y}) =0.2460
Pr(x) = 0.5625	Pr(~y) = 0.4375	Pr($K \cup J$ = {x})= 0.2460
Pr(~x) = 0.4375	Pr(~y) = 0.4375	Pr($K \cup J$ = {Ø}) =0.1914

Event Probabilities for Plan $L \cup J$

x-Events	y-Events	z-Events	Joint Event
Pr(x) = 0.5625	Pr(y) = 0.25	Pr(z) = 0.25	Pr($L \cup J$ ={x,y,z}) = 0.0351
Pr(~x) = 0.4375	Pr(y) = 0.25	Pr(z) = 0.25	Pr($L \cup J$ ={y,z}) =0.0273
Pr(x) = 0.5625	Pr(~y) = 0.75	Pr(z) = 0.25	Pr($L \cup J$ ={x,z})= 0.1054
Pr(~x) = 0.4375	Pr(~y) = 0.75	Pr(z) = 0.25	Pr($L \cup J$ = {z}) =0.0820
Pr(x) = 0.5625	Pr(y) = 0.25	Pr(~z) = 0.75	Pr($L \cup J$ = {x,y}) = 0.1054
Pr(~x) = 0.4375	Pr(y) = 0.25	Pr(~z) = 0.75	Pr($L \cup J$ = {y}) =0.0820
Pr(x) = 0.5625	Pr(~y) = 0.75	Pr(~z) = 0.75	Pr($L \cup J$ = {x})= 0.3154
Pr(~x) = 0.4375	Pr(~y) = 0.75	Pr(~z) = 0.75	Pr($L \cup J$ = {Ø}) =0.2460

Hunter (2004) shows that there exists a monotone one-to-one mapping from a set of unique genes in a species collection and the Blackwell (1953) information index when genes are "information" in the sense defined by Weitzman (1998) (i.e. independent Bernoulli trials). Let $B:|S| \rightarrow \Theta$, where Θ is the set of admissable Blackwell indecies and S is a set whose members are unique "genes." For concreteness, refer to specific values of the index which correspond to the *ex post* (i.e. after the extinction event) cardinality of a conservation plan as:

$$\theta = \begin{cases} \theta_\phi, \text{ when } |\mathrm{S}| = 0 \\ \theta_1, \text{ when } |\mathrm{S}| = 1 \\ \theta_2, \text{ when } |\mathrm{S}| = 2 \\ \theta_3, \text{ when } |\mathrm{S}| = 3 \end{cases}$$

Then, the probability of each of the possible outcomes for the feasible conservation plans can be used find the likelihood that a conservation plan will actually preserve an information set of a given ranking, in the Blackwell sense. For each of the plans above, the marginal distribution of possible *ex post* Blackwell rankings is given in the following tables:

Marginal distributions for feasible plans

Plan $K \cup J$	Plan $L \cup J$	Plan $L \cup K$
$\Pr(\theta_\phi) = 0.1914$	$\Pr(\theta_\phi) = 0.2460$	$\Pr(\theta_\phi) = 0.2460$
$\Pr(\theta_1) = 0.4920$	$\Pr(\theta_1) = 0.4794$	$\Pr(\theta_1) = 0.4794$
$\Pr(\theta_2) = 0.3164$	$\Pr(\theta_2) = 0.2381$	$\Pr(\theta_2) = 0.2381$
$\Pr(\theta_3) = 0$	$\Pr(\theta_3) = 0.0351$	$\Pr(\theta_3) = 0.0351$

Radner and Stiglitz (1980) have shown that the *value* of information is an increasing *often convex* function of the Blackwell measure of information. Then, a comparison of the prizes (i.e., information sets) in the "species lotteries" (i.e., conservation plans) can be accomplished using the cumulative distribution functions for each feasible plan (Rothschild and Stiglitiz, 1971). A robust comparison requires that the CDF's be ranked using first-order stochastic dominance because of the increasing returns to information's value noted by Radner and Stiglitz(1980). In the table below, the CDF's for the two informationally distinct conservation plans is given:

Cumulative distributions for distinct plans

Plan $K \cup J$	Plan $L \cup J$
$\Pr(\theta \le \theta_\phi) = 0.1914$	$\Pr(\theta \le \theta_\phi) = 0.2460$
$\Pr(\theta \le \theta_1) = 0.6834$	$\Pr(\theta \le \theta_1) = 0.7254$
$\Pr(\theta \le \theta_2) = 1$	$\Pr(\theta \le \theta_2) = 0.9635$
$\Pr(\theta \le \theta_3) = 1$	$\Pr(\theta \le \theta_3) = 1$

Because neither of theses CDF's lies everywhere below (weakly) the other, first-order dominance is not established. Thus, without specifying details regarding the planner's preferences, it cannot be claimed that either of these plans will definitively generate higher expected utility. Clearly, the example could easily be modified to show cases where first-order dominance is established in the reserve site selection problem. However, in most cases first-order dominance (FSD) would be the basis for a partial ordering of reserve sites, as the present example illustrates.

Concluding Remarks

The example presented illustrates many of the most important points we have mentioned in exposing our approach toward welfare comparisons. Chief among these is that ordering the set of social states is not trivial. Taking into consideration the relative value of different policy outcomes due to the many use and non-use values for collections of species can be daunting. That being said, once the ranking of the social state has been determined, comparing the CDFs induced by various policy proposal to see if any them stochastically dominates the others is relatively straightforward. This approach promises to identify any policies that are cost-effective in a robust sense, as well as those that are optimal only for a particular set of welfare weights.

References

Armsworth, P.R., B.E. Kendall, and F.W. Davis. 2004. An introduction to biodiversity concepts for environmental economists. *Resource and Energy Economics* **26**: 115-136.

Arthur, J.L., J.D. Camm, R.G. Haight, C.A. Montegomery, and S. Polasky. 2004. Weighing conservation objectives: maximum expected coverage versus endangered species protection. *Ecological Applications* **14**(6): 1936-1945.

Camm, J.D., S. Polasky, A. Solow, and B. Csuti. 1996. A note on optimal algorithms for reserve site selection. *Biological Conservation* **78**: 353-355.

Camm, J.D., S.K. Norman, S. Polasky, and A.R. Solow. 2002. Nature reserve site selection to maximize expected species covered. *Operations Research* **50**(6): 946-955.

Church, R.L., D.M. Stoms, and F.W. Davis. 1996. Reserve selection as a maximal covering location problem. *Biological Conservation* **76**: 105-112.

Davis, F.W., C. Costello, and D. Stoms. 2006. Efficient conservation in a utility-maximization framework. *Ecology and Society* **11**(1):33. [online] URL:http://www.ecologyandsociety.org/vol11/iss1/art33/

Margules, C.R., and R.L. Pressey. 2000. Systematic conservation planning. *Nature* **405**: 243-253.

Metrick, A., and M.L. Weitzman. 1998. Conflicts and choices in biodiversity preservation. *Journal of Economic Perspectives* **12**(2): 21-34.

Nicholson, E., and H.P. Possingham. 2005. Objectives for multiple-species conservation planning. *Conservation Biology* **20**(3): 871-881.

Polasky, S., C. Costello, and A. Solow. 2005. The economics of biodiversity. *Handbook of Natural Resource Economics Volume 5 ???*
Purvis, A., J.L. Gittleman, G. Cowlishaw, and G.M. Mace. 2000. Predicting extinction risk in declining species. *Proceedings of the Royal Society of London B* **267**: 1947-1952.

Rothschild, M. and Stiglitz J.E. 1970. Increasing Risk I: A Definition. *Journal of Economic Theory* **3**: 65-73.

S. Polasky, Camm, J.D., A. Solow, B. Csuti, D. White, and R. Ding. 2000. Choosing reserve networks with incomplete species information. *Biological Conservation* **94**: 1-10.

Williams, J.C. 2006. Optimal reserve site selection with distance requirements. *Computers & Operations Research.* [in press].

Williams, P.H., and M.B. Araújo. 2000. Using probabilithy of persistence to identify important areas for biodiversity conservation. *Proceedings of the Royal Society of London B* **267**: 1959-1966.

Wilson, K.A., Westphal, M.I., H.P. Possingham, and J. Elith. 2005. Sensitivity of conservation planning to different approaches to using predicted species distribution data. *Biological Conservation* **122**: 99-112.

Blackwell, D., Equivalent Comparison of Experiments. 1953. *Annals of Mathematics and Statistics*, 24: 265-72.

Hunter, G., "Noah's Non-concavity: On the existence of non-trivial interior solutions to the problem of cost-effective conservation planning," mimeo California State University, 2004.

Radner, R. and Stiglitz, J., "A Nonconcavity in the value of information," in Marcel Boyer and Richard Kihlstrom (eds.), *Bayesian Models in Economic Theory*, Elsevier Science Publications, 1984, pp. 33-52.

Rothschild, M. and Stiglitz, J. 1970. Increasing Risk: 1. A Definition. *Journal of Economic Theory **2(3)***: 225-243.

Weitzman, M., "The Noah's Ark Problem," *Econometrica*, 66 (1998), 1279-1298.

Green, J and J.-J. Laffont. (1979). *Incentives in Public Decision Making.* Amsterdam:North-Holland.