

On the Foundations of *Expected Expected* Utility

Craig Boutilier

Department of Computer Science
University of Toronto
Toronto, ON, M5S 3H5, CANADA
cebly@cs.toronto.edu

Abstract

Intelligent agents often need to assess user utility functions in order to make decisions on their behalf, or predict their behavior. When uncertainty exists over the precise nature of this utility function, one can model this uncertainty using a distribution over utility functions. This view lies at the core of games with incomplete information and, more recently, several proposals for incremental preference elicitation. In such cases, decisions (or predicted behavior) are based on computing the *expected expected* utility (EEU) of decisions with respect to the distribution over utility functions. Unfortunately, decisions made under EEU are sensitive to the precise representation of the utility function. We examine the conditions under which EEU provides for sensible decisions by appeal to the foundational axioms of decision theory. We also discuss the impact these conditions have on the enterprise of preference elicitation more broadly.

1 Introduction

Most work on the foundations of decision theory—specifically on the justification of expected utility—has focused on *personal decision making*, that is, settings where a decision is being made by the “holder” of the utility function. Of course the decision maker may not be fully aware of (or have fully articulated) her utility function. The process of articulation is complex, and much work in decision analysis deals with preference elicitation and decision framing to help the decision maker formulate her decision problem [11]. However, this work is primarily concerned with eliciting enough information about preference tradeoffs to allow an (approximately) optimal decision to be made. While an analyst can never be sure about the true nature of the decision maker’s utility function, this uncertainty is not generally characterized explicitly, though its impact is often minimized through sensitivity analysis and related techniques.

Recent emphasis has been placed on the development of automated decision tools, where a decision is being made on behalf of a user whose utility function is imprecisely known. As in goal programming or other forms of interactive optimization, a space of possible utility functions is

usually maintained (often by imposing constraints on trade-off weights). A decision can be made based on this set of feasible utility functions. For example, Pareto optimal decisions can be identified [21; 18], or models based on minimax regret can be used to choose a specific decision [11; 2; 20]. In each of these models, the uncertainty regarding the utility function is characterized by the feasible utility set.

Somewhat less common is work in which the system’s uncertainty about a user’s utility function is quantified probabilistically. Some recent examples include [5; 6; 1]. In this work, a distribution over utility functions is assumed. The expected utility of a decision is determined not just by taking expectation over the outcomes of that decision, but also expectation over the space of possible utility functions. We use the term *expected expected utility (EEU)* to denote the value of a decision computed in this way. Elicitation strategies can be informed using the current distribution over utility functions. For example, value of information can be used to determine whether the improvement in decision quality given by a piece of information outweighs the cost of obtaining that information. Thus, characterizing one’s uncertainty over possible utility functions in a probabilistic fashion, and using EEU to determine decision quality, has much to recommend it from the point of view of elicitation.

Decision making using distributions over utility functions has been considered in other contexts. For example, Cyert and de Groot consider problems in sequential decision making in which uncertainty in the underlying utility function is represented probabilistically [8; 9]. Fishburn [10] also addresses this problem (as we discuss below). Harsanyi’s formulation of games with incomplete information as Bayesian games [12; 13] relies critically on distributions over payoff functions, and virtually the entire literature on in this area adopts this perspective [7; 15].¹

In all of this work, the EEU concept is used to determine the value of decisions in the context of an uncertain utility function. Unfortunately, while EEU has an intuitive appeal, this scheme is sensitive to positive affine transformations of

¹In some sense, much work in collaborative filtering [3; 16] and related models [4] can be viewed as incorporating distributions over utility functions. However, these are used for purposes of classification (i.e., determining a unique utility function for a particular user) and generally uncertainty in utility is not accounted for when making decisions.

the utility functions in question. Implicit in such a scheme is a *commensurability assumption* that allows the quantities present in the different utility functions to be meaningfully compared and combined. This is not always the case. The aim of this paper is to describe certain conditions under which this commensurability assumption can be justified by appeal to the foundational axioms for decision theory as proposed by von Neumann and Morgenstern [19] and Savage [17].

The setting we consider is one in which an agent for a decision maker or user is uncertain about the user's preferences, but wishes to recommend (or take) decisions on the user's behalf. Fishburn [10] has considered the problems of the foundations of expected expected utility from a somewhat different perspective. He considers the problem in which a decision maker is uncertain about the set of consequences she might face and considers combining utility functions over different consequence sets. Unfortunately, his results cannot be applied (except in a trivial way) to the situation above.²

We begin by defining the problem of decision making given uncertainty over utility functions and the EEU concept. We then examine the sensitivity of EEU to the precise representation of the underlying utility functions, and propose an interpretation of utility uncertainty that allows one to prescribe "canonical" utility function representations under specific circumstances. We conclude with a brief discussion of the implications these considerations have for "practical" elicitation.

2 Expected Expected Utility

We begin by establishing notation and basic background with a quick overview of expected utility and then define the notion of expected expected utility formally.

2.1 Expected Utility

Assume a *decision scenario* consisting of a finite set of possible *decisions* D , a finite set of possible outcomes (or states) S , and a distribution function $\Pr_d \in \Delta(S)$, for each $d \in D$. The term $\Pr_d(s)$ denotes the probability of outcome s being realized if the system takes decision d . $\Pr_d(s)$ can be viewed as a vector p_d whose i th component is $\Pr_d(s_i)$ (given a suitable enumeration of outcomes in S).

A *utility function* $u : S \rightarrow \mathbb{R}$ associates utility $u(s)$ with each outcome s . We will generally view u as a $|S|$ -dimensional vector $\mathbf{u}$ whose i th component u_i is $u(s_i)$.³ The *expected utility* of decision d w.r.t. u is:

$$EU(d, u) = p_d \mathbf{u} = \sum_{i \in S} \Pr_d(s_i) u_i.$$

The optimal decision d^* w.r.t. u is that with *maximum expected utility* (MEU).

²While our results are general, it is unclear how profitable it is to model a decision maker's uncertainty about her *own* utility function. It can be argued that such uncertainty should be viewed as "traditional" uncertainty about future outcomes, context, etc. Rather than take a stand on this issue, we simply emphasize that an agent can be genuinely uncertain about a user's utility function, and that our model and results apply (in a practical way) to such a setting.

³If u is represented using some more concise model, $\mathbf{u}$ is simply the vector of parameters required for that model.

It is well known that utility functions are invariant under positive affine transformations. That is, the relative expected utility of any pair of decisions (in any decision scenario) will be unaltered by such a transformation of a utility function. This implies that the optimal decision in any decision scenario is unaffected by such a transformation.

More precisely, von Neumann and Morgenstern equate [19] (classes of) utility functions with preferences over *lotteries*. Let $\langle p_1, s_1; p_2, s_2; \dots p_n, s_n \rangle$, where $\sum_i p_i = 1$, denote a *simple lottery* over outcomes, with each outcome s_i obtained with probability p_i . As a shorthand, we sometimes omit outcomes whose probability is zero. Note that an outcome s_i is itself a (trivial) simple lottery, and that each decision d induces a simple lottery $l(d)$ over outcomes. A *compound lottery* is a lottery whose elements may be further lotteries. Let $\succ$ a preference function over lotteries, with $l_1 \succ l_2$ meaning that l_1 is strictly preferred to l_2 . The relations $\succeq$, $\prec$, $\preceq$, and $\sim$ are defined in the usual way. Assuming certain (relatively uncontroversial) axioms restricting the form of the preference function $\succ$, von Neumann and Morgenstern show that there exists a utility function $u_\succ$ that exactly represents $\succ$ in the following sense: $EU(d, u_\succ) > EU(d', u_\succ)$ iff $l(d) \succ l(d')$. Furthermore, the utility function $u_\succ$ is unique up to positive affine transformation. Because of this, we can partition the space of utility functions into equivalence classes, each corresponding to a unique preference ordering $\succ$. We denote this class by $[\succ]$.

Several of the axioms used to equate utility functions with preferences over lotteries are listed here. These are based on the work of von Neumann and Morgenstern [19] and Savage [17], though the particular form used here is drawn from [14]. We use $s_\top$ to denote some most preferred outcome in S (i.e., $s_\top \succeq s, \forall s \in S$), and $s_\perp$ to denote some least preferred outcome (these are guaranteed to exist due to other axioms).

Monotonicity $\langle p, s_\top; 1 - p, s_\perp \rangle \succeq \langle q, s_\top; 1 - q, s_\perp \rangle$ iff $p \geq q$.

Continuity For each s_i , there is some p such that $s_i \sim \langle p, s_\top; 1 - p, s_\perp \rangle$.

Reduction of Compound Lotteries Let l be the lottery $\langle p_1, l_1; \dots; p_n, l_n \rangle$ where each l_i is a lottery of the form $\langle q_1^i, l_1^i; \dots; q_m^i(i), l_m^i(i) \rangle$. Let $\tilde{l}_k, k \leq K$ denote the K (unique) lotteries within the set $\{l_j^i : i \leq n, j \leq m(i)\}$. Let r be the (reduced) lottery over the (unique) component lotteries $\tilde{l}_k$ with probability $\tilde{p}_k = \sum \{p_i q_j^i : l_j^i = \tilde{l}_k\}$ associated with each $\tilde{l}_k$. Then $r \sim l$.

2.2 Uncertainty over Utility Functions

An agent will often not know the user's utility function u with certainty. We model this uncertainty using a density P over the set of utility functions $U \subseteq \mathbb{R}^{|S|}$ (or a distribution over a finite support set contained in U). If a system makes a decision d under such conditions of uncertainty, the expected utility of d must reflect this. We consider the following definition for the expected utility of d given density P over U :

$$EU(d, P) = \int p_d \mathbf{u} P(\mathbf{u}) d\mathbf{u}.$$

We refer to this as the *expected expected utility* (EEU) of decision d , since it is the expectation of $EU(d, u)$ w.r.t. $P(U)$.

This definition is precisely that used in [5; 6; 1] in the context of utility elicitation, and also that used in much other work involving uncertainty over utility [12; 8; 9]. In such a state of uncertainty—or *belief state*—the optimal decision is that d^* with maximum EEU $EU(d^*, P)$. We denote by $EU(P)$ the value of being in belief state P , assuming one is forced to make a decision:

$$EU(P) = \max_{d \in D} EU(d, P).$$

We call this generic decision rule *the MEEU decision rule*, by analogy with the classical MEU decision rule.

EEU seems to be a fairly natural concept given probabilistically quantified uncertainty over utilities. The fact that it occurs in many different contexts certainly attests to this fact. Unfortunately, the proposed definition can induce certain anomalies, as we examine below.

3 Justifying MEEU

3.1 Loss of Invariance

The results of von Neumann and Morgenstern suggest that the decisions one makes with respect to belief state P over U should be invariant to legitimate transformations of the elements of U . Certainly, this would be a desirable feature of the MEEU decision rule. One might even claim that the decision rule can only be considered useful if it satisfies this condition. In general, unfortunately, this is not the case.

As a simple illustration, suppose we have a domain with two outcomes s_1 and s_2 , and a distribution P that assigns probability 0.5 to $u_1 = \langle 1, 3 \rangle$ and probability 0.5 to $u_2 = \langle 2, 1 \rangle$. Suppose we use the MEEU decision rule in this context, by computing

$$EU(d, P) = \sum_{u_i} \sum_{s_j} p_d(s_j) u_i(s_j) P(u_i).$$

and choosing the decision d^* with maximum expected utility $EU(d^*, P)$. Then a decision that accords higher probability to s_2 will be preferred to one that gives lower probability to s_2 . However, if we transform u_2 into $u'_2 = \langle 20, 10 \rangle$, the relative utilities of these decisions will be reversed. Thus, the MEEU decision rule is not insensitive to transformations of individual utility functions with positive support. Note that we are not suggesting that agent's will arbitrarily transform some utility functions and not others.⁴ Rather, the question is: which representation of a specific utility function (e.g., u_2 in the example) should be adopted in the first place?

One possible way to deal with this problem is to recognize that a utility function is simply a convenient (and nonunique) way of expressing preferences over lotteries. Rather than working with utility functions, we could work explicitly with a density over preference functions (in fact, we will do this implicitly below). Unfortunately, the set of lotteries over which a preference ordering is defined is uncountable; therefore, some compact representation (of the *individual* preference functions) is needed. But this is precisely the role of

⁴If the same transformation is applied to all functions with positive support, the MEEU decision is unchanged.

a utility function—to serve as a concise representation of a preference function over lotteries.

This gives rise to the question of how to choose a representative utility function from each equivalence class $[\succ]$ that allows formal justification of the MEEU decision rule, and under what circumstances such representatives exist.

3.2 A Lottery Interpretation of MEU

We give a formal justification for the MEEU rule under a specific condition: *we assume the existence of a known best and worst outcome*. That is, each utility function with positive support has the same best outcome $s_\top$ and worst outcome $s_\perp$. We also insist that the user is not indifferent to these alternatives, that is, that $s_\top$ must be strictly preferred to $s_\perp$.⁵ We call such utility functions *extremum equivalent*. In many settings, such as those involving active preference elicitation, restricting attention to a set of extremum equivalent utility functions is not problematic. One simply needs to ask the user to identify her most and least preferred outcomes (these need not be unique, but only one such representative need be identified).

Once we know outcomes $s_\top$ and $s_\perp$, we insist that our belief state P assign nonzero measure only to normalized utility functions u in which $u(s_\top) = c_\top$ and $u(s_\perp) = c_\perp$ for some constants $c_\top > c_\perp$. For convenience, we will assume $c_\top = 1$ and $c_\perp = 0$, but nothing crucial depends on these choices. Under these conditions, every preference ordering $\succ$ for lotteries has a unique utility function representation satisfying the normalization constraints.

We note that each preference ordering corresponds to a unique set of indifference conditions between outcomes and standard lotteries. A *standard lottery* is one of the form $\langle p, s_\top; 1 - p, s_\perp \rangle$. By the Continuity Axiom, we have, for each outcome s_i and each preference ordering $\succ$,

$$s_i \sim \langle p_i^\succ, s_\top; 1 - p_i^\succ, s_\perp \rangle,$$

for some unique $p_i^\succ$. We denote by $l(\succ, s_i)$ this standard lottery. This implies that $u_\succ(s_i) = p_i^\succ$ in the normalized utility function $u_\succ$ corresponding to $\succ$.⁶

Now given a belief state P , how does one compare the value of different decisions d ? This can be reduced to a question of how one compares the value of guaranteed outcomes s_i , since decisions are just lotteries over outcomes. Given a particular preference ordering $\succ$, we have $s_i \sim l(\succ, s_i)$. However, the choice facing the agent for the decision maker involves uncertainty over the true preference ordering. Thus, we can view the outcome s_i as a compound lottery, where first a preference ordering is chosen according to distribution P , and then the standard gamble involving s_i is played. This will induce a new preference ordering $\succ_P$ over outcomes (and decisions). We illustrate this for the case of a discrete distribution over two preference orderings $\succ$ and $\succ'$, with probabilities q and $1 - q$ respectively. In this case, the outcome s_i is equated with the compound lottery

$$s_i \sim \langle q, l(\succ, s_i); 1 - q, l(\succ', s_i) \rangle.$$

⁵If $s_\top \sim s_\perp$, then the decision problem is trivial since each decision is equally preferred. Indeed, in some axiomatizations, nontriviality is imposed on legitimate preferences [11].

⁶If the normalizing constants differ from 0/1, then utility is some linear function of $p_i^\succ$. This has no impact on our argument.

By reduction of compound lotteries, this means that

$$s_i \sim_P \langle q \cdot p_i^{\succ} + (1-q) \cdot p_i^{\succ'}, s_{\top}; q \cdot (1-p_i^{\succ}) + (1-q) \cdot (1-p_i^{\succ'}), s_{\perp} \rangle.$$

Similarly, for any other outcome s_j , we have

$$s_j \sim_P \langle q \cdot p_j^{\succ} + (1-q) \cdot p_j^{\succ'}, s_{\top}; q \cdot (1-p_j^{\succ}) + (1-q) \cdot (1-p_j^{\succ'}), s_{\perp} \rangle.$$

By the Monotonicity Axiom, we have $s_i \succ_P s_j$ iff

$$q \cdot p_i^{\succ} + (1-q) \cdot p_i^{\succ'} > q \cdot p_j^{\succ} + (1-q) \cdot p_j^{\succ'}.$$

Since $p_i^{\succ} = u_{\succ}(s_i)$ (and similarly for the other terms), we have $s_i \succ_P s_j$ iff $EU(s_i, P) > EU(s_j, P)$ (where we emphasize that here we are treating outcomes as deterministic decisions that guarantee the corresponding outcomes).⁷ From this one can easily show that for any two decisions d, d' (that induce distributions over outcomes), $l(d) \succ_P l(d')$ iff $EU(d, P) > EU(d', P)$.

This argument applies to arbitrary discrete distributions, and can be generalized to continuous densities as follows:

$$\begin{aligned} s_i \succ_P s_j & \\ \text{iff } \left\langle \int_U l(\succ_u, s_i) \right\rangle \succ_P \left\langle \int_U l(\succ_u, s_j) \right\rangle & \\ \text{iff } \left\langle \int_U p_i^{\succ u} P(u), s_{\top}; 1 - \int_U p_i^{\succ u} P(u), s_{\perp} \right\rangle \succ_P & \\ \left\langle \int_U p_j^{\succ u} P(u), s_{\top}; 1 - \int_U p_j^{\succ u} P(u), s_{\perp} \right\rangle & \\ \text{iff } \int_U p_i^{\succ u} P(u) > \int_U p_j^{\succ u} P(u) & \\ \text{iff } \int_U u(s_i) P(u) > \int_U u(s_j) P(u) & \\ \text{iff } EU(s_i, P) > EU(s_j, P) & \end{aligned}$$

Here the first step refers to a compound lottery over an continuous set of component (simple) lotteries, while we assume in second step that a such a compound lottery can be reduced to a simple lottery in an analogous way to the reduction of a finite compounded lottery.

Thus under the assumption that one can identify a best and worst outcome, the MEEU decision rule can be justified for use with normalized (extremum equivalent) utility functions by appeal to the foundational axioms of decision theory, and an interpretation of uncertainty over utility as a lottery over the lotteries defined by the component utility functions.

We now formalize the legitimacy of EEU and MEE.

Definition 1 Let $\{\succ_i\}$ be a set of extremum equivalent preference relations with respect to finite outcome set S , with best and worst outcomes $s_{\top}$ and $s_{\perp}$, respectively. Let P be a density over $\{\succ_i\}$. For any $\succ_i$, let $p_s^{\succ_i}$ denote the tradeoff probability for state s in its standard gamble w.r.t. $\succ_i$; that is,

$$s \sim \langle p_s^{\succ_i}, s_{\top}; 1 - p_s^{\succ_i}, s_{\perp} \rangle.$$

⁷Note that if we chose other normalizing constants for our best and worst outcomes, then we would have that $u_{\succ}(s_i)$ is some linear function of $p_i^{\succ}$; but this linear function is identical for each term in the equation, so the conclusion holds.

The *aggregate standard gamble* for $s \in S$ induced by $\{\succ_i\}$ and P is defined:

$$\left\langle \int_{\succ_i} p_s^{\succ_i} P(\succ_i), s_{\top}; 1 - \int_{\succ_i} p_s^{\succ_i} P(\succ_i), s_{\perp} \right\rangle.$$

Let l be any lottery w.r.t. S . The *aggregate reduction* of l is the standard gamble

$$R_P(l) = \langle p_P^l, s_{\top}; 1 - p_P^l, s_{\perp} \rangle$$

obtained by replacing every outcome s in l by its aggregate standard gamble and reducing it to a standard gamble in the usual fashion. The *aggregate preference relation* $\succ$ (w.r.t. S) induced by $\{\succ_i\}$ and P is given by

$$l_1 \succ l_2 \text{ iff } p_P^{l_1} > p_P^{l_2}.$$

Theorem 1 Let $\{\succ_i\}$ and P be defined as above. Let $\{u_i\}$ be a set of utility functions (one for each $\succ_i$, consistent with $\succ_i$) such that $u_i(s_{\top}) = c_{\top}$ and $u_i(s_{\perp}) = c_{\perp}$, for all u_i and two fixed constants $c_{\top} > c_{\perp} \geq 0$. Then the utility function $u(s) = \int_{\succ_i} u_i(s) P(\succ_i)$ is consistent with the aggregate preference relation $\succ$ induced by $\{\succ_i\}$ and P .

Extremum equivalence is thus sufficient to ensure commensurability, as it puts all utility functions on a common scale. It is important to realize that the scale dictated by the best and worst outcomes *cannot vary*, since these are truly best and worst outcomes; we return to this point below. It appears to be much more difficult to apply this type of argument to densities over utility functions that are not extremum equivalent.

Fishburn [10] considers the problem of EEU when a decision maker is uncertain about the nature of the consequence sets she will face. He proposes foundational axioms that justify the use EEU to compare decisions. However, the setting is rather different: specifically, Fishburn requires that any consequences that two utility functions have in common be ranked identically. In our context, where each utility function lies over the same consequence set, the Fishburn axioms impose overly stringent requirements. It is interesting to note that Fishburn requires something akin to extremum equivalence, namely, that there exist two consequences common to the domains of each utility function such that one of the consequences is preferred to the other in each function.

4 Dealing with Small Worlds

It is important to realize that the best and worst outcomes with which one calibrates must either be truly best and worst outcomes from the decision maker's standpoint, or they themselves must be calibrated. Using Savage's [17] terminology, we must be careful to distinguish "small worlds" reasoning from "grand worlds." Consider the case where the set of outcomes is restricted to the subset of outcomes that are possible given the set of actions in a specific decision scenario. But assume there exist outcomes outside the domain of the restricted scenario for which the user has concrete preferences. Let's refer to the set of restricted outcomes as *local*, while the space of all outcomes is *global*.

We might imagine determining the best and worst *local* outcomes, $s_\top^l$ and $s_\perp^l$, respectively, and engaging in the elicitation process using standard gambles with respect to these extreme local outcomes. Unfortunately, this is not sufficient to justify the MEEU decision rule. The difficulty is that the user's degrees of preference for $s_\top^l$ and $s_\perp^l$ may themselves be characterized by uncertainty. Specifically, $s_\top^l$ can be equated with a standard gamble $\langle p_\top^l, s_\top; 1 - p_\top^l, s_\perp \rangle$ (where $s_\top, s_\perp$ denote the globally best and worst outcomes). Similarly, $s_\perp^l \sim \langle p_\perp^l, s_\top; 1 - p_\perp^l, s_\perp \rangle$. But the relevant gamble probabilities— $p_\top^l$ and $p_\perp^l$ —may vary with the utility function and will not be known if attention is restricted to the *small worlds* domain. As a consequence, our distribution over *local* utility functions is insufficient to justify MEEU.

To illustrate, suppose that we believe the user's utility function is either u_1 (with probability q_1) or u_2 (with probability $q_2 = 1 - q_1$). Furthermore, suppose that we consider only the projection of these onto the set of local outcomes, calibrating two outcomes s_i and s_j with respect to "local" standard gambles. Suppose that we have the following tradeoffs with respect to the local extrema in u_1 :

$$s_i \sim_1 \langle p_i^1, s_\top^l; (1 - p_i^1), s_\perp^l \rangle$$

$$s_j \sim_1 \langle p_j^1, s_\top^l; (1 - p_j^1), s_\perp^l \rangle$$

and similarly in u_2 :

$$s_i \sim_2 \langle p_i^2, s_\top^l; (1 - p_i^2), s_\perp^l \rangle$$

$$s_j \sim_2 \langle p_j^2, s_\top^l; (1 - p_j^2), s_\perp^l \rangle.$$

Using EEU, we have

$$s_i \sim \langle q_1 p_i^1 + q_2 p_i^2, s_\top^l; q_1(1 - p_i^1) + q_2(1 - p_i^2), s_\perp^l \rangle$$

$$s_j \sim \langle q_1 p_j^1 + q_2 p_j^2, s_\top^l; q_1(1 - p_j^1) + q_2(1 - p_j^2), s_\perp^l \rangle.$$

Now, were we to place this *small world* in the *grand world* context, we might say that the local best and worst outcomes, $s_\top^l$ and $s_\perp^l$, have specific lottery probabilities with respect to the global utility functions. Let's assume that we have

$$s_\top^l \sim_1 \langle p_\top^1, s_\top; (1 - p_\top^1), s_\perp \rangle$$

$$s_\perp^l \sim_1 \langle p_\perp^1, s_\top; (1 - p_\perp^1), s_\perp \rangle$$

$$s_\top^l \sim_2 \langle p_\top^2, s_\top; (1 - p_\top^2), s_\perp \rangle$$

$$s_\perp^l \sim_2 \langle p_\perp^2, s_\top; (1 - p_\perp^2), s_\perp \rangle.$$

In other words, in (the normalized, global counterpart of) utility function u_1 , we have $u_1(s_\top^l) = p_\top^1$ and $u_1(s_\perp^l) = p_\perp^1$; while in u_2 we have $u_2(s_\top^l) = p_\top^2$ and $u_2(s_\perp^l) = p_\perp^2$. Applying EEU and reduction of compound lotteries, we see that

$s_i \sim$

$$\begin{aligned} & \langle p_\top^1 q_1 p_i^1 + p_\top^2 q_2 p_i^2 + p_\perp^1 q_1(1 - p_i^1) + p_\perp^2 q_2(1 - p_i^2), s_\top; \\ & (1 - p_\top^1) q_1 p_i^1 + (1 - p_\top^2) q_2 p_i^2 \\ & + (1 - p_\perp^1) q_1(1 - p_i^1) + (1 - p_\perp^2) q_2(1 - p_i^2), s_\perp \rangle \end{aligned}$$

and

$s_j \sim$

$$\begin{aligned} & \langle p_\top^1 q_1 p_j^1 + p_\top^2 q_2 p_j^2 + p_\perp^1 q_1(1 - p_j^1) + p_\perp^2 q_2(1 - p_j^2), s_\top; \\ & (1 - p_\top^1) q_1 p_j^1 + (1 - p_\top^2) q_2 p_j^2 \\ & + (1 - p_\perp^1) q_1(1 - p_j^1) + (1 - p_\perp^2) q_2(1 - p_j^2), s_\perp \rangle. \end{aligned}$$

Hence, $s_i \succ s_j$ iff

$$\begin{aligned} & p_\top^1 q_1 p_i^1 + p_\top^2 q_2 p_i^2 + p_\perp^1 q_1(1 - p_i^1) + p_\perp^2 q_2(1 - p_i^2) \\ & > p_\top^1 q_1 p_j^1 + p_\top^2 q_2 p_j^2 + p_\perp^1 q_1(1 - p_j^1) + p_\perp^2 q_2(1 - p_j^2). \end{aligned}$$

Unfortunately, this condition is not equivalent to the MEEU rule in the original small worlds domain. Specifically, this condition cannot generally be assessed without having some assessment of the global tradeoff probabilities $p_\top^1$, etc. In other words, to accurately compare two small world outcomes given uncertainty about the *local* utility functions, we have to explicitly assess our uncertainty about the *range* of values the local extrema can take with respect to the *global* utility function. Thus while one can generally use small world reasoning in the classic decision-theoretic setting, its use is problematic in the EEU framework.

Fortunately, the MEEU principle can be recovered if we make one simple assumption: that the global tradeoffs associated with the best and worst local outcomes do not vary with the utility function (at least, with those utility functions having positive support). For example, to continue with the illustration above, suppose that the tradeoff probabilities associated with the two utilities u_1 and u_2 are identical; that is, $p_\top^1 = p_\top^2$ and $p_\perp^1 = p_\perp^2$. (Whether the probabilities are known is irrelevant, all that matters is that they are known to be identical.) If this is so we have $s_i \succ s_j$ iff

$$\begin{aligned} & p_\top^1(q_1 p_i^1 + q_2 p_i^2) + p_\perp^1(q_1(1 - p_i^1) + q_2(1 - p_i^2)) \\ & > p_\top^1(q_1 p_j^1 + q_2 p_j^2) + p_\perp^1(q_1(1 - p_j^1) + q_2(1 - p_j^2)) \end{aligned}$$

Since the inner terms in each expression sum to one, and $p_\top^1 > p_\perp^1$ (since $s_\top^l \succ s_\perp^l$), we must have $s_i \succ s_j$ iff

$$q_1 p_i^1 + q_2 p_i^2 > q_1 p_j^1 + q_2 p_j^2$$

This, of course, is the expression for MEEU.

Indeed, we don't have to assume that the global tradeoff probabilities are fixed. It is sufficient to assume that the unknown range of (global) utility spanned by $s_\top^l$ and $s_\perp^l$ is probabilistically independent of the unknown local utility function u . This too is sufficient to allow MEEU to be used justifiably. We omit the argument, which is straightforward.

5 Consequences for Preference Elicitation

The considerations above have implications for practical preference elicitation. From a foundational perspective, calibration of utilities relative to known best and worst outcomes is required if decisions are to be based on EEU. In the case of incremental elicitation, where EEU is used to determine value of information, we must first obtain a prior over a set of extremum equivalent utility functions before engaging in such deliberations. Fortunately, it often seems to make sense to determine best and worst outcomes beforehand, and engage in "serious" elicitation after this initial calibration.

Another important question: how does one determine priors over utility functions? Utility function databases [4] could be used. This poses some problems regarding interpersonal utility comparison for which there are no especially compelling solutions. When using EEU, things are a bit worse: we need to construct priors conditioned on the observed or

elicited best and worst outcomes. Given a prior over arbitrary utility functions, as long as decisions using EEU are not made until the determination of best and worst outcomes is completed, this poses no difficulties. An alternative, in certain scenarios, is to suppose that certain outcomes are universally most and least preferred (e.g., in medical contexts, death can be used as the latter). This may be hard to justify formally, but from a practical point of view will be quite useful and (one hopes) have little impact on actual decision quality.

The issue of small worlds also poses certain problems. From the point of view of practical elicitation, the prospect of calibrating some small set of relevant outcomes using a user's "global" utility function is unappealing. Fortunately, as argued above, for a given individual, strength of preference can be assumed fixed for the best and worst outcomes, which allows things to carry through. Strength of preference may prove to be important however when trading off the increase in EEU with the effort associated with the elicitation process. Furthermore, this can have an important impact on the construction of priors from databases of utility functions.

6 Concluding Remarks

Decision making when the underlying utility function is unknown is an important problem in game theory, interactive optimization, and preference elicitation. Quantifying this uncertainty using distributions over utility functions has a number of appealing qualities, and quite naturally leads to the notion of *expected* expected utility, a concept that has been used in several different lines of research.

The aim of this paper is to point out that expectations taken with respect to utility function distributions require some care. More precisely, the operation of expected expected utility only makes sense (from a foundational standpoint) when the distribution is over extremal equivalent utility functions. While this has certain implications for practical utility elicitation, we have argued that this requirement is not overly stringent from a practical perspective.

Acknowledgements

Thanks to Radford Neal whose comments on preference elicitation led to this investigation and to Moshe Tennenholtz for suggesting interesting connections. This work was supported by the Institute for Robotics and Intelligent Systems (IRIS), and the Natural Sciences and Engineering Research Council of Canada (NSERC).

References

- [1] Craig Boutilier. A POMDP formulation of preference elicitation problems. In *Proc. Eighteenth National Conference on Artificial Intelligence*, pages 239–246, Edmonton, 2002.
- [2] Craig Boutilier, Fahiem Bacchus, and Ronen I. Brafman. UCP-Networks: A directed graphical representation of conditional utilities. In *Proc. Seventeenth Conference on Uncertainty in Artificial Intelligence*, pages 56–64, Seattle, 2001.
- [3] Jack S. Breese, David Heckerman, and Carl Kadie. Empirical analysis of predictive algorithms for collaborative filtering. In *Proc. Fourteenth Conference on Uncertainty in Artificial Intelligence*, pages 43–52, Madison, WI, 1998.
- [4] U. Chajewska, L. Getoor, J. Norman, and Y. Shahar. Utility elicitation as a classification problem. In *Proc. Fourteenth Conference on Uncertainty in Artificial Intelligence*, pages 79–88, Madison, WI, 1998.
- [5] Urszula Chajewska and Daphne Koller. Utilities as random variables: Density estimation and structure discovery. In *Proc. Sixteenth Conference on Uncertainty in Artificial Intelligence*, pages 63–71, Stanford, 2000.
- [6] Urszula Chajewska, Daphne Koller, and Ronald Parr. Making rational decisions using adaptive utility elicitation. In *Proc. Seventeenth National Conference on Artificial Intelligence*, pages 363–369, Austin, TX, 2000.
- [7] Tommy Chin-Chiu Chan and Sérgio Ribeiro da Costa Werlang. The Bayesian foundations of the solution concepts of games. *Journal of Economic Theory*, 45:370–391, 1988.
- [8] Richard M. Cyert and Morris H. de Groot. Adaptive utility. In Maurice Allais and Ole Hagen, editors, *Expected Utility Hypothesis and the Allais Paradox*, pages 223–241. D. Reidel, Dordrecht, 1979.
- [9] Morris H. de Groot. Decision making with an uncertain utility function. In Bernt P. Stigum and Fred Wenstøp, editors, *Foundations of Utility and Risk Theory with Applications*, pages 371–384. D. Reidel, Dordrecht, 1983.
- [10] Peter C. Fishburn. *The Foundations of Expected Utility*. D. Reidel, Dordrecht, 1982.
- [11] Simon French. *Decision Theory*. Halsted Press, New York, 1986.
- [12] John C. Harsanyi. Games with incomplete information played by Bayesian players. part I: The basic model. *Management Science*, 14:159–182, 1967.
- [13] John C. Harsanyi. Games with incomplete information played by Bayesian players. part II: Bayesian equilibrium points. *Management Science*, 14:320–334, 1968.
- [14] R. D. Luce and H. Raiffa. *Games and Decisions*. Wiley, New York, 1957.
- [15] Roger B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, 1991.
- [16] David M. Pennock and Eric Horvitz. Collaborative filtering by personality diagnosis: A hybrid memory- and model-based approach. In *Proc. Sixteenth Conference on Uncertainty in Artificial Intelligence*, pages 473–480, Stanford, 2000.
- [17] Leonard J. Savage. *The Foundations of Statistics*. Wiley, New York, 1954.
- [18] Edward A. Sykes and Chelsea C. White, III. Multiobjective intelligent computer-aided design. *IEEE Transactions on Systems, Man and Cybernetics*, 21(6):1498–1511, 1991.
- [19] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, 1944.
- [20] Tianhan Wang and Craig Boutilier. Incremental utility elicitation with the minimax regret decision criterion. In *Proc. Eighteenth International Joint Conference on Artificial Intelligence*, this volume, Acapulco, 2003.
- [21] Chelsea C. White, III, Andrew P. Sage, and Shigeru Dozono. A model of multiattribute decisionmaking and trade-off weight determination under uncertainty. *IEEE Transactions on Systems, Man and Cybernetics*, 14(2):223–229, 1984.