Great Expectations.
Part II: Generalized Expected Utility as a Universal Decision Rule*

Francis C. Chu and Joseph Y. Halpern
Department of Computer Science
Cornell University
Ithaca, NY 14853, U.S.A.
Email: {fcc,halpern}@cs.cornell.edu

Abstract

Many different rules for decision making have been introduced in the literature. We show that a notion of generalized expected utility proposed in a companion paper [Chu and Halpern, 2003] is a universal decision rule, in the sense that it can represent essentially all other decision rules.

1 Introduction

A great deal of effort has been devoted to studying decision making. A standard formalization decribes the choices a decision maker (DM) faces as acts, where an *act* is a function from states to consequences. Many decision rules (that is, rules for choosing among acts, based on the tastes and beliefs of the DM) have been proposed in the literature. Some are meant to describe how "rational" agents should make decisions, while others aim at modeling how real agents actually make decisions. Perhaps the best-known approach is that of *maximizing expected utility* (EU). Normative arguments due to Savage [1954] suggest that rational agents should behave as if their tastes are represented by a real-valued utility function on the consequences, their beliefs about the likelihood of events (i.e., sets of states) are represented by a probability measure, and they are maximizing the expected utility of acts with respect to this utility and probability.

Despite these normative arguments, it is well known that EU often does not describe how people actually behave when they make decisions [Resnik, 1987]; thus EU is of limited utility if we want to model (and perhaps predict) how people will behave. As a result, many alternatives to EU have been proposed in the literature (see, for example, [Gul, 1991; Gilboa and Schmeidler, 1989; Giang and Shenoy, 2001; Quiggin, 1993; Schmeidler, 1989; Yaari, 1987]). Some of these rules involve representations of beliefs by means other than a (single) probability measure; in some cases, beliefs and tastes are combined in ways other than the standard way which produces expected utility; yet other cases, such as Maximin and Minimax Regret [Resnik, 1987], do not require a representation of beliefs at all.

*Work supported in part by NSF under grants IIS-0090145 and CTC-0208535 and by the DoD Multidisciplinary University Research Initiative (MURI) program administered by ONR under grant N00014-01-1-0795.

In a companion paper [Chu and Halpern, 2003], we propose a general framework in which to study and compare decision rules. The idea is to define a generalized notion of expected utility (GEU), where a DM's beliefs are represented by plausibility measures [Friedman and Halpern, 1995] and the DM's tastes are represented by general (i.e., not necessarily real-valued) utility functions. We show there that every preference relation on acts has a GEU representation. Here we show that GEU is universal in a much stronger sense: we show that essentially all decision rules have GEU representations. The notion of representing one decision rule using another seems to be novel. Intuitively, decision rules are functions from tastes (and beliefs) to preference relations, so a representation of a decision rule is a representation of *a function,* not a preference relation.

Roughly speaking, given two decision rules R_1 and R_2, an R_1 *representation of* 'R_2 is a function T that maps inputs of R_2 to inputs of R_1 that contain the same representation of tastes (and beliefs) such that $\mathcal{R}_1(\tau(x)) = \mathcal{R}_2(x)$. Thus, T models, in a precise sense, a user of R_2 as a user of R_1, since r preserves tastes (and beliefs). We show that a large collection of decision rules have GEU representations and characterize the collection. Essentially, a decision rule has a GEU representation iff it is *uniform* in a precise sense. It turns out that there are well-known decision rules, such as maximizing Choquet expected utility (CEU) [Schmeidler, 1989] (which essentially assumes that the DM is representing beliefs using a Dempster-Shafer belief function Bel, and then maximizing CEU with respect to Bel), that have no GEU representations. This is because r is not allowed to modify the representation of the tastes (and beliefs). We then define a notion *of ordinal representation,* in which r is allowed to modify the representation of the tastes (and beliefs), and is required to preserve only the ordinal aspect of the tastes (and beliefs). We show that almost all decision rules, including CEU, have ordinal GEU representations.

There seems to be no prior work in the literature that considers how one decision rule can represent another. Perhaps the closest results to our own are those of Lehmann [2001]. He proposes a "unified general theory of decision" that contains both quantitative and qualitative decision theories. He considers a particular decision rule he calls *Expected Qualitative Utility Maximization,* which allows utilities to be nonstandard real numbers; he defines a certain preorder on the nonstandard reals and makes decisions based on maximizing expected utility (with respect to that preorder). That his framework has EU as a special case is immediate, since for the standard reals, his preorder reduces to the standard order on the reals. He argues informally that Maximin is a special case of his approach, so that his approach can capture aspects of more qualitative decision making as well. It is easy to see that Lehmann's approach is a special case of GEU; his rule is clearly not universal in our sense.

2 Preliminaries

To make this paper self-contained, much of the material in the first three subsections of this section is taken (almost verbatim) from [Chu and llalpern, 2003].

2.1 Plausibility, Utility, and Expectation Domains

Since one of the goals of this paper is to provide a general framework for all of decision theory, we want to represent the tastes and beliefs of the DMs in as general a framework as possible. To this end, we use plausibility measures to represent the beliefs of the DMs and (generalized) utility functions to represent their tastes.

A *plausibility domain* is a set P, partially ordered by $\preceq_P$ (so $\preceq_P$ is a reflexive, antisymmetric, and transitive relation), with two special elements $\bot_P$ and $\top_P$, such that for all $x \in P$, $\bot_P \preceq_P x \preceq_P \top_P$. A function $\mathrm{Pl} : 2^S \to P$ is a *plausibility measure* iff

Pl1. $\mathrm{Pl}(\emptyset) = \bot$,

Pl2. $\mathrm{Pl}(S) = \top$, and

Pl3. if $X \subseteq Y$ then $\mathrm{Pl}(X) \preceq \mathrm{Pl}(Y)$.

As pointed out in [Friedman and Halpern, 1995], plausibility measures generalize, not only probability, but a host of other representations of uncertainty as well. A *utility domain* is a set U endowed with a reflexive binary relation $\precsim_U$. Intuitively, elements of U represent the strength of likes and dislikes of the DM while elements of P represent the strength of her beliefs.

Once we have plausibility and utility, we want to combine them to form expected utility. To do this, we introduce expectation domains, which have utility domains, plausibility domains, and operators $\oplus$ (the analogue of $+$) and $\otimes$ (the analogue of $\times$).[1] More formally, an *expectation domain* is a tuple $E = (U, P, V, \otimes, \oplus)$, where $(U, \precsim_U)$ is a utility domain, $(P, \preceq_P)$ is a plausibility domain, $(V, \precsim_V)$ is a valuation domain (where $\precsim_V$ is a reflexive binary relation), $\otimes : P \times U \to V$, and $\oplus : V \times V \to V$. (As usual, we omit subscripts when they are clear.) We have four requirements on expectation domains:

*E*1. $(x \oplus y) \oplus z = x \oplus (y \oplus z)$;

*E*2. $x \oplus y = y \oplus x$;

*E*3. $\top \otimes x = x$;

*E*4. $(U, \precsim_U)$ is a substructure of $(V, \precsim_V)$.

[1] Sometimes we use x to denote Cartesian product; the context will always make it clear whether this is the case.

E_1 and E_2 say that 0 is associative and commutative. £3 says that T is the left-identity of ® and *E4* ensures that the expectation domain respects the relation on utility values.

The *standard expectation domain,* which we denote E, is (R, [0,1], R, +, x), where the ordering on each domain is the standard order on the reals.

2.2 Decision Situations and Decision Problems

A *decision situation* describes the objective part of the circumstance that the DM faces (i.e., the part that is independent of the tastes and beliefs of the DM). Formally, a decision situation is a tuple *A* — *{A,* 5, C), where

- *S* is the set of states of the world,
- *C* is the set of consequences, and
- *A* is a set of acts (i.e., a set of functions from *S* to *C).*

An act *a* is *simple* iff its range is finite. That is, *a* is simple if it has only finitely many consequences. Many works in the literature focus on simple acts (e.g., [Fishburn, 1987]). We assume in this paper that *A* contains only simple acts; this means that we can define (generalized) expectation using finite sums, so we do not have to introduce infinite series or integration for arbitrary expectation domains. Note that all acts are guaranteed to be simple if either *S* or *C* is finite, although we do not assume that here.

A decision problem is essentially a decision situation together with information about the tastes (and beliefs) of the DM; that is, a decision problem is a decision situation together with the subjective part of the circumstance that faces the DM. Formally, a *nonplausibilistic decision problem* is a tuple *(A, U,* u), where

- *A* = (A,5, *C)* is a decision situation,
- *U* is a utility domain, and
- u : *C* —> *U* is a utility function.

A *plausibilistic decision problem* is a tuple *(A, E,* u, Pl), where

- *A* — (A, 5, *C)* is a decision situation,
- E= (U,, *P, V,* ®, ®) is an expectation domain,
- u : *C* —> *U* is a utility function, and
- Pl : 2^{S} —> *P* is a plausibility measure.

We could have let a plausibilistic decision problem be simply a nonplausibilistic decision problem together with a plausibility domain and a plausibility measure, without including the other components of expectation domains. However, this turns out to complicate the presentation (see below).

We say that *V* is *standard* iff its utility domain is R (and, if D is plausibilistic, its plausibility measure is a probability measure and its expectation domain is E).

2.3 Expected Utility

Let *V* be a decision problem with *S* as the set of states, *U* as the utility domain, and u as the utility function. Each act *a* of *V* induces a *utility random variable* u_a : *S* —► *U* as follows: $u_a(s)$ = u(a(s)). If in addition *V* is plausibilistic with *P* as the plausibility domain and Pl as the plausibility measure, then each a also induces a *utility lottery* $\ell_a^{\mathrm{Pl},\mathbf{u}} : \mathrm{ran}(\mathbf{u}_a) \rightarrow P$ as follows: $\ell_a^{\mathrm{Pl},\mathbf{u}}(u) = \mathrm{Pl}(\mathbf{u}_a^{-1}(u))$. If $\mathcal{D}$ is in fact standard (so $E = \mathbb{E}$ and Pl is a probability measure Pr), we can identify the expected utility of act a with the expected value of $\mathbf{u}_a$ with respect to Pr, computed in the standard way:

$$\mathbf{E}_{\mathrm{Pr}}(\mathbf{u}_a) = \sum_{x \in \mathrm{ran}(\mathbf{u}_a)} \mathrm{Pr}(\mathbf{u}_a^{-1}(x)) \times x. \qquad (2.1)$$

We can generalize (2.1) to an arbitrary expectation domain $E = (U, P, V, \otimes, \oplus)$ by replacing $+$, $\times$, and Pr by $\oplus$, $\otimes$, and Pl, respectively. This gives us

$$\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a) = \bigoplus_{x \in \mathrm{ran}(\mathbf{u}_a)} \mathrm{Pl}(\mathbf{u}_a^{-1}(x)) \otimes x. \qquad (2.2)$$

We call (2.2) the *generalized EU* (GEU) of act a. Clearly (2.1) is a special case of (2.2).

2.4 Decision Rules

Intuitively, a decision rule tells the DM what to do when facing a decision problem in order to get a preference relation on acts—e.g., compare the expected utility of acts. Just as we have nonplausibilistic decision problems and plausibilistic decision problems, we have nonplausibilistic decision rules and plausibilistic decision rules. As the name suggests, (non)plausibilistic decision rules are defined on (non)plausibilistic decision problems.

We do not require decision rules to be defined on all decision problems. For example, (standard) EU is defined only on standard plausibilistic decision problems. More formally, a *(non)plausibilistic decision rule* R is a function whose domain, denoted dom(R.), is a subcollection of the collection of (non)plausibilistic decision problems, and whose range, denoted ran (R), is a collection of preference relations on acts. If $\mathcal{D} \in$ doin(R) and a_1 and a_2 are acts in V, then we write

$$a_1 \precsim_{\mathcal{R}(\mathcal{D})} a_2 \text{ iff } (a_1, a_2) \in \mathcal{R}(\mathcal{D}).$$

Here are a few examples of decision rules:

- GEU is a plausibilistic decision rule whose domain consists of all plausibilistic decision problems. Given a decision p lem $\mathcal{D} = (\mathcal{A}, E, \ldots)$, where $E = (U, P, V, \oplus, \otimes)$, have $a_1 \precsim_{\mathrm{GEU}(\mathcal{D})} a_2$ iff $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{a_1}) \precsim_V \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{a_2})$ for all acts a_1, a_2 of $\mathcal{A}$. Note that GEU would not be a decision rule according to this definition if plausibilistic decision problems contained only a utility function and a plausibility measure, and did not include the other components of expectation domains.

- Of course, standard EU is a decision rule (whose domain consists of all standard plausibilistic decision problems).

- Maximin is a nonplausibilistic decision rule that orders acts according to their worst-case consequence. It is a conservative rule; the "best" act according to Maximin is the one with the best worst-case consequence. Intuitively, Maximin views Nature as an adversary that always pick a state that realizes the worst-case consequence, no matter what act the DM chooses. The domain of (standard) Maximin consists of nonplausibilistic decision problems with real-valued utilities. Given an act a and a real-valued utility function $\mathbf{u}$, let $\mathbf{w}_{\mathbf{u}}(a) = \min_{s \in S} \mathbf{u}_a(s)$. Then given a decision problem $\mathcal{D} = (\mathcal{A}, \mathbb{R}, \mathbf{u})$, $a_1 \precsim_{\mathrm{Maximin}(\mathcal{D})} a_2$ iff $\mathbf{w}_{\mathbf{u}}(a_1) \leq \mathbf{w}_{\mathbf{u}}(a_2)$.

- Minimax Regret (REG) is based on a different philosophy. It tries to hedge a DM's bets, by doing reasonably well no matter what the actual state is. It is also a nonplausibilistic rule. As a first step to defining it, given a nonplausibilistic decision problem $\mathcal{D} = ((A, S, C), \mathbb{R}, \mathbf{u})$, for each state $s \in S$, let $\overline{u}_s = \sup_{a \in A} \mathbf{u}_a(s)$; that is, $\overline{u}_s$ is the least upper bound of the utilities in state s. The *regret* of a in state .s, denoted $\mathbf{r}(a, s)$, is $\overline{u}_s - \mathbf{u}_a(s)$; note that no act can do better than a by more than r(a, s) in state s. Let $\overline{\mathbf{r}}(a) = \sup_{s \in S} \mathbf{r}(a, s)$. For example, suppose that $\overline{\mathbf{r}}(a) = 2$ and the DM picks a. Suppose that the DM then learns that the true state is s_o and is offered a chance to change her mind. No matter what act she picks, the utility of the new act cannot be more than 2 higher then $\mathbf{u}_a(s_0)$. REG orders acts by their regret and thus takes the "best" act to be the one that minimizes $\overline{\mathbf{r}}(a)$. Intuitively, this rule tries to minimize the regret that a DM would feel if she discovered what the situation actually was: the "I wish 1 had done a_2 instead of a_1" feeling. Thus, $a_1 \precsim_{\mathrm{REG}(\mathcal{D})} a_2$ iff $\overline{\mathbf{r}}(a_1) \geq \overline{\mathbf{r}}(a_2)$. Like Maximin, Nature is viewed as an adversary that would pick a state that maximizes regret, no matter what act the DM chooses. It is well known that, in general, Maximin, REG, and EU give different recommendations [Resnik, 1987].

- The Maxmin Expected Utility rule (MMEU) [Gilboa and Schmeidler, 1989] assumes that a DM's beliefs are represented by a set V of probability measures. Act a_1 is preferred to a_2 if the worst-case expected utility of a_1 (taken over all the probability measures in V) is at least as large as the worst-case expected utility of a_2. Thus MMEU is, in a sense, a hybrid of EU and Maximin. To view MMEU as a function on decision problems, we must first show how to represent a set of probability measures as a single plausibility measure. We do this using an approach due to Halpern [2001]. Let the plausibility domain $P = [0, 1]^{\mathcal{P}}$, that is, all functions from V to [0,1], ordered pointwise; in other words, $p \preceq_P q$ iff $p(\mathrm{Pr}) \leq q(\mathrm{Pr})$ for all $\mathrm{Pr} \in \mathcal{P}$. Thus, in this domain, $\perp$ is the constant function 0 and T is the constant function 1. For each $X \subseteq S$, let $f_X \in P$ be the function that evaluates each probability measure in V at A"; that is, $f_X(\mathrm{Pr}) = \mathrm{Pr}(X)$ for all $\mathrm{Pr} \in \mathcal{P}$. Let $\mathrm{Pl}_{\mathcal{P}}(X) = f_X$; it is easy to verify that Plp is a plausibility measure. We view Plp as a representation of the set V of probability measures; clearly V can be recovered from $\mathrm{Pl}_{\mathcal{P}}$. The domain of MMEU consists of all plausibilistic decision problems of the form $\mathcal{D} = ((A, S, C), (\mathbb{R}, [0,1]^{\mathcal{P}}, V, \oplus, \otimes), \mathbf{u}, \mathrm{Pl}_{\mathcal{P}})$, where V is a set of probability measures on 2^S, and $a_1 \precsim_{\mathrm{MMEU}(\mathcal{D})} a_2$ iff $\inf_{\mathrm{Pr} \in \mathcal{P}} \mathbf{E}_{\mathrm{Pr}}(\mathbf{u}_{a_1}) \leq \inf_{\mathrm{Pr} \in \mathcal{P}} \mathbf{E}_{\mathrm{Pr}}(\mathbf{u}_{a_2})$. Note that this definition ignores $\oplus$, $\otimes$, and V.

- (Dempster-Shafer) *belief functions* [Dempster, 1967] are a representation of uncertainty that generalize probability. That is, every probability measure is a belief function, but the converse is not necessarily true.[2] Given a belief function Bel, it is well-known that there exists a set $\mathcal{P}_{\mathrm{Bel}}$ of probability measures such that for all $X \subseteq S$, Bel(A^r) = $\inf_{\mathrm{Pr} \in \mathcal{P}_{\mathrm{Bel}}} \mathrm{Pr}(X)$ [Dempster, 1967]. A notion of expected

[2] Due to lack of space, we assume that the reader is familiar with belief functions.

utility for belief functions was defined by Choquet [1953] as follows. Given an act a and a real-valued utility function $\mathbf{u}$, if $\mathrm{ran}(\mathbf{u}_a) = \{u_1, \ldots, u_n\}$ and $u_1 < \cdots < u_n$, then

$$\mathbf{E}_{\mathrm{Bel}}(\mathbf{u}_a) = u_1 + \sum_{i=2}^{n} \mathrm{Bel}(X_i) \times (u_i - u_{i-1}), \qquad (2.3)$$

where $X_i = \mathbf{u}_a^{-1}(\{u_i, \ldots, u_n\})$. It is easy to check (2.3) agrees with (2.1) if Bel is a probability measure. Moreover,

$$\mathbf{E}_{\mathrm{Bel}}(\mathbf{u}_a) = \inf_{\Pr \in \mathcal{P}_{\mathrm{Bel}}} \mathbf{E}_{\Pr}(\mathbf{u}_a) \qquad (2.4)$$

in general. The Choquet expected utility (CEU) rule has as its domain decision problems of the form $\mathcal{D} = (\mathcal{A}, \mathbb{E}, \mathbf{u}, \mathrm{Bel})$. Define $a_1 \precsim_{\mathrm{CEU}(\mathcal{D})} a_2$ iff $\mathbf{E}_{\mathrm{Bel}}(\mathbf{u}_{a_1}) \leq \mathbf{E}_{\mathrm{Bel}}(\mathbf{u}_{a_2})$. It is immediate from (2.4) that if $\mathcal{D}_{\mathcal{P}_{\mathrm{Bel}}}$ is the decision problem that results from $\mathcal{D}$ by replacing Bel by $\mathrm{Pl}_{\mathcal{P}_{\mathrm{Bel}}}$, and replacing the plausibility domain $[0,1]$ in the expectation domain by $[0,1]^{\mathcal{P}_{\mathrm{Bel}}}$, then $a_1 \precsim_{\mathrm{CEU}(\mathcal{D})} a_2$ iff $a_1 \precsim_{\mathrm{MMEU}(\mathcal{D}_{\mathcal{P}_{\mathrm{Bel}}})} a_2$.

3 Representing Decision Rules

Given a decision rule $\mathcal{R}$ and a preference relation $\precsim_A$ on the set of acts A, an $\mathcal{R}$ *representation of* $\precsim_A$ is basically a decision problem $\mathcal{D} \in \mathrm{dom}(\mathcal{R})$ such that $\mathcal{R}(\mathcal{D}) = \precsim_A$ (and the set of acts in $\mathcal{D}$ is A). In other words, an $\mathcal{R}$ representation of $\precsim_A$ makes $\mathcal{R}$ relate acts in A the way $\precsim_A$ relates them, so we can model a DM whose preference relation is $\precsim_A$ as a user of $\mathcal{R}$. In [Chu and Halpern, 2003] we show that every preference relation on acts has a GEU representation. Here we want to extend the notion of representation to decision rules; intuitively, we want an $\mathcal{R}_1$ representation of $\mathcal{R}_2$ to allow us to model a user of $\mathcal{R}_2$ as a user of $\mathcal{R}_1$. To make this precise, we need a few definitions.

Two (plausibilistic) decision problems $\mathcal{D}_1$ and $\mathcal{D}_2$ are *congruent*, denoted $\mathcal{D}_1 \cong \mathcal{D}_2$, iff they involve the same decision situation, utility domain, and utility function (and, if both are plausibilistic, the same plausibility domain and plausibility measure as well). Note that if $\mathcal{D}_1 \cong \mathcal{D}_2$, then they agree on the tastes (and beliefs) of the DM, so if they are both nonplausibilistic, then $\mathcal{D}_1 = \mathcal{D}_2$, and if they are both plausibilistic, then they differ only in the V, $\precsim_V$, $\oplus$, and $\otimes$ components of their expectation domains.

A *decision rule transformation* τ is a function that maps inputs of one decision rule $\mathcal{R}_2$ to the inputs of another rule $\mathcal{R}_1$. A decision rule transformation τ is an $\mathcal{R}_1$ *representation of* $\mathcal{R}_2$ iff $\mathrm{dom}(\tau) = \mathrm{dom}(\mathcal{R}_2)$ and for all $\mathcal{D} \in \mathrm{dom}(\mathcal{R}_2)$,

$$\tau(\mathcal{D}) \cong \mathcal{D} \text{ and } \mathcal{R}_1(\tau(\mathcal{D})) = \mathcal{R}_2(\mathcal{D}).$$

Thus a DM that uses $\mathcal{R}_2$ to relate acts based on her tastes (and beliefs) behaves as if she is using $\mathcal{R}_1$, since $\tau(\mathcal{D}) \cong \mathcal{D}$ and $\mathcal{R}_1(\tau(\mathcal{D})) = \mathcal{R}_2(\mathcal{D})$.

Note that $\tau(\mathcal{D}) = \mathcal{D}$ is a GEU representation of EU. We now consider some less trivial examples.

Example 3.1: To see that Maximin has a GEU representation, let $E_{\max} = (\mathbb{R}, \{0,1\}, \mathbb{R}, \min, \times)$, and let $\mathrm{Pl}_{\max}$ be the plausibility measure such that $\mathrm{Pl}_{\max}(U)$ is 0 if $X = \emptyset$ and 1 otherwise. If $\mathcal{D} = (\mathcal{A}, \mathbb{R}, \mathbf{u})$, where $\mathcal{A} = (A, S, C)$, then it is easy to check that $\mathbf{E}_{\mathrm{Pl}_{\max}, E_{\max}}(\mathbf{u}_a) = \mathbf{w}_{\mathbf{u}}(a)$. Take $\tau(\mathcal{D}) = (\mathcal{A}, E_{\max}, \mathbf{u}, \mathrm{Pl}_{\max})$. Clearly $\tau(\mathcal{D}) \cong \mathcal{D}$: the utility function has not changed. Moreover, it is immediate that $\mathrm{GEU}(\tau(\mathcal{D})) = \mathrm{Maximin}(\mathcal{D})$. ∎

Example 3.2: To see that Minimax Regret (REG) has a GEU representation, for ease of exposition, we take dom(REG) to consist of standard decision problems $\mathcal{D} = ((A, S, C), \mathbb{R}, \mathbf{u})$ such that $M_{\mathcal{D}} = \sup_{a \in A, s \in S} \mathbf{u}_a(s) < \infty$. (If $M_{\mathcal{D}} = \infty$, given the restriction to simple acts, it is easy to show that all acts have infinite regret.) Let $E_{\mathrm{reg}} = (\mathbb{R}, [0,1], \mathbb{R}, \min, \otimes)$, where $x \otimes y = y - \log(x)$ if $x > 0$, and $x \otimes y = 0$ if $x = 0$. Note that $\bot = 0$ and $\top = 1$. Clearly, min is associative and commutative, and $\top \otimes r = r - \log(1) = r$ for all $r \in \mathbb{R}$. Thus, E_{reg} is an expectation domain.

For $\emptyset \neq X \subseteq S$, define $M_X = \sup_{a \in A, s \in X} \mathbf{u}_a(s)$. Note that $M_S = M_{\mathcal{D}} < \infty$; also if $X \subseteq Y$, then $M_X \leq M_Y$. Let $\mathrm{Pl}_{\mathcal{D}}(\emptyset) = 0$ and $\mathrm{Pl}_{\mathcal{D}}(X) = e^{M_X - M_S}$. It is easy to verify that $\mathrm{Pl}_{\mathcal{D}}$ is a plausibility measure.

It is easy to check that $\mathbf{E}_{\mathrm{Pl}_{\mathcal{D}}, E_{\mathrm{reg}}}(\mathbf{u}_a) = M_{\mathcal{D}} - \mathbf{r}_{\mathbf{u}}(a)$ for all acts $a \in A$. Let $\tau(\mathcal{D}) = (\mathcal{A}, E_{\mathrm{reg}}, \mathbf{u}, \mathrm{Pl}_{\mathcal{D}})$. Since the utility function has not changed, $\tau(\mathcal{D}) \cong \mathcal{D}$; furthermore, higher expected utility corresponds to lower regret, so $\mathrm{GEU}(\tau(\mathcal{D})) = \mathrm{REG}(\mathcal{D})$. ∎

Example 3.3: To see that MMEU has a GEU representation, let $\mathcal{D} = (\mathcal{A}, (\mathbb{R}, [0,1]^{\mathcal{P}}, \widehat{V}, \widehat{\oplus}, \widehat{\otimes}), \mathbf{u}, \mathrm{Pl}_{\mathcal{P}}) \in \mathrm{dom}(\mathrm{MMEU})$. Let $E_{\mathcal{P}} = (\mathbb{R}, [0,1]^{\mathcal{P}}, \mathbb{R}^{\mathcal{P}}, \oplus, \otimes)$, where $\oplus$ is pointwise function addition, $\otimes$ is scalar multiplication, and

$$f \precsim_{\mathbb{R}^{\mathcal{P}}} g \text{ iff } \inf_{\Pr \in \mathcal{P}} f(\Pr) \leq \inf_{\Pr \in \mathcal{P}} g(\Pr).$$

Note that we can identify $\mathbb{R}$ with the constant functions in $\mathbb{R}^{\mathcal{P}}$, so $\mathbb{R}$ can be viewed as a substructure of $\mathbb{R}^{\mathcal{P}}$. With these definitions, $E_{\mathcal{P}}$ is an expectation domain.

Let $\tau(\mathcal{D}) = (\mathcal{A}, E_{\mathcal{P}}, \mathbf{u}, \mathrm{Pl}_{\mathcal{P}})$. It is immediate from the definition of $\precsim_{\mathbb{R}^{\mathcal{P}}}$ that

$$a \precsim_{\mathrm{GEU}(\tau(\mathcal{D}))} b \text{ iff } \inf_{\Pr \in \mathcal{P}} \mathbf{E}_{\Pr}(\mathbf{u}_a) \leq \inf_{\Pr \in \mathcal{P}} \mathbf{E}_{\Pr}(\mathbf{u}_b).$$

Thus $\mathrm{GEU}(\tau(\mathcal{D})) = \mathrm{MMEU}(\mathcal{D})$; furthermore, it is clear that $\tau(\mathcal{D}) \cong \mathcal{D}$, since the utility function and plausibility measure have not changed. ∎

Although it can represent many decision rules, GEU cannot represent CEU. We can in fact characterize the conditions under which a decision rule is representable by GEU.

There is a trivial condition that a decision rule must satisfy in order for it to have a GEU representation. Intuitively, a decision rule R respects utility if R relates acts of constant utility according to the relation between utility values. Formally, a decision rule 7v *respects utility* iff for all $\mathcal{D} \in$ dom(R.) with A as the set of acts, S as the set of states, U as the utility domain, and u as the utility function, for all $a_1, a_2 \in A$, if $\mathbf{u}_{a_i}(s) = u_i$ for all states $s \in S$, then

$$a_1 \precsim_{\mathcal{R}(\mathcal{D})} a_2 \text{ iff } u_1 \precsim_U u_2. \qquad (3.5)$$

We say that R *weakly respects utility* iff (3.5) holds for all constant acts (but not necessarily for all acts of constant utility). It is easy to see that GEU respects utility, since $\top \otimes u = u$ for all $u \in U$ and $(U, \precsim_U)$ is a substructure of $(V, \precsim_V)$. Thus if R does not respect utility, it has no GEU representation. While respecting utility is a necessary condition for a decision rule to have a GEU representation, it is not sufficient. It is also necessary for the decision rule to treat acts that behave in similar ways similarly.

Two acts a_1, a_2 in a decision problem $\mathcal{D}$ are *indistinguishable*, denoted $a_1 \sim_{\mathcal{D}} a_2$ iff either

- $\mathcal{D}$ is nonplausibilistic and $\mathbf{u}_{a_1} = \mathbf{u}_{a_2}$, or
- $\mathcal{D}$ is plausibilistic, and $\ell_{a_1}^{\mathrm{Pl},\mathbf{u}} = \ell_{a_2}^{\mathrm{Pl},\mathbf{u}}$,

where $\mathbf{u}$ is the utility function of $\mathcal{D}$ and Pl is the plausibility measure of $\mathcal{D}$. In the nonplausibilistic case, two acts are indistinguishable if they induce the same utility random variable; in the plausibilistic case, they are indistinguishable if they induce the same utility lottery.

A decision rule $\mathcal{R}$ is uniform if it respects indistinguishability. More formally, $\mathcal{R}$ is *uniform* iff for all $\mathcal{D} \in \mathrm{dom}(\mathcal{R})$ and a_1, a_2, b_1, b_2 acts of $\mathcal{D}$ such that $a_i \sim_{\mathcal{D}} b_i$,

$$a_1 \precsim_{\mathcal{R}(\mathcal{D})} a_2 \text{ iff } b_1 \precsim_{\mathcal{R}(\mathcal{D})} b_2.$$

Intuitively, we can think of utility random variables and utility lotteries as descriptions of what an act *a* does in terms of the tastes (and beliefs) of the DM. If 7v is uniform, we can view 7v as relating the acts indirectly by relating their descriptions.

As the following theorem shows, all uniform decision rules have GEU representations.

Theorem 3.4: *For all decision rules R,1Z has a GEU representation iff R, is uniform and R respects utility.*

Proof: The " i f direction is somewhat similar in spirit to the proof of Theorem 3.6, given below; due to the lack of space, we omit this direction.

For the "only if" direction, suppose that τ is a GEU representation of $\mathcal{R}$ and let $\mathcal{D}_0 \subset \mathrm{dom}(\mathcal{R})$ be arbitrary. Suppose that a_1, a_2, b_1, b_2 are acts of $\mathcal{D}_0$ such that $a_i \sim_{\mathcal{D}_0} b_i$. It is easy to check that if $\mathcal{D} = (A, E, \mathbf{u}, \mathrm{Pl}) \cong \mathcal{D}_0$, then $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{a_i}) = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{b_i})$. Thus for all plausibilistic $\mathcal{D}$, if $\mathcal{D} \cong \mathcal{D}_0$, then $a_1 \precsim_{\mathrm{GEU}(\mathcal{D})} a_2$ iff $b_1 \precsim_{\mathrm{GEU}(\mathcal{D})} b_2$. Since τ is a GEU representation of $\mathcal{R}$, $\tau(\mathcal{D}_0) \cong \mathcal{D}_0$ and $\mathcal{R}(\mathcal{D}_0) = \mathrm{GEU}(\tau(\mathcal{D}_0))$. It follows then that $a_1 \precsim_{\mathcal{R}(\mathcal{D}_0)} a_2$ iff $b_1 \precsim_{\mathcal{R}(\mathcal{D}_0)} b_2$; thus $\mathcal{R}$ is uniform.

Now suppose that a_1 and a_2 are two acts of constant utility, say u_1 and u_2, respectively, of $\mathcal{D}_0$. Since $\tau(\mathcal{D}_0) \cong \mathcal{D}_0$, a_i is still an act of constant utility u_i in $\tau(\mathcal{D}_0)$. Note that $a_1 \precsim_{\mathcal{R}(\mathcal{D}_0)} a_2$ iff $a_1 \precsim_{\mathrm{GEU}(\tau(\mathcal{D}_0))} a_2$ iff $u_1 \precsim_U u_2$, where U is the utility domain of $\mathcal{D}_0$, since τ is a GEU representation of $\mathcal{R}$. Thus $\mathcal{R}$ respects utility. ∎

Most of the decision rules we have discussed are uniform. However, CEU is not, as the following example shows:

Example 3.5: Let $\mathcal{D}_* = ((A, S, C), \mathrm{E}, \mathbf{u}, \mathrm{Bel})$, where

- $A = \{a_1, a_2\}$; $S = \{s_1, s_2, s_3\}$; $C = \{1, 2, 3\}$;
- $\mathbf{u}(j) = j$, for $j = 1, 2, 3$;
- $a_1(s_j) = j$ and $a_2(s_j) = 3 - j$, for $j = 1, 2, 3$; and
- Bel is the belief function such that $\mathrm{Bel}(X) = 1$ if $\{s_1, s_2\} \subseteq X$ and $\mathrm{Bel}(X) = 0$ otherwise.

Since $\mathbf{u}_{a_i}^{-1}(j)$ is a singleton, $\mathrm{Bel}(\mathbf{u}_{a_i}^{-1}(j)) = 0$ for $i = 1, 2$ and $j = 1, 2, 3$; thus $a_1 \sim_{\mathcal{D}_*} a_2$. On the other hand, by definition,

$$\mathbf{E}_{\mathrm{Bel}}(\mathbf{u}_{a_1}) = 1 + \mathrm{Bel}(X_{23})(2 - 1) + \mathrm{Bel}(X_3)(3 - 2) = 1,$$

while

$$\mathbf{E}_{\mathrm{Bel}}(\mathbf{u}_{a_2}) = 1 + \mathrm{Bel}(X_{12})(2 - 1) + \mathrm{Bel}(X_1)(3 - 2) = 2,$$

where $X_{ij} = \{s_i, s_j\}$ and $X_k = \{s_k\}$. It follows that CEU is not uniform, and so has no GEU representation. ∎

The reader may have noticed an incongruity here. Example 3.3 shows that MMEU has a GEU representation; moreover, as shown earlier, MMEU produces essentially the same order on acts as CEU. However, CEU has no GEU representation. There is no contradiction to Theorem 3.4 here: There is no decision problem $\mathcal{D}$ such that $\mathcal{D} \cong \mathcal{D}_*$ (from Example 3.5) and $\mathrm{GEU}(\mathcal{D}) = \mathrm{CEU}(\mathcal{D}_*)$. However, $\mathrm{GEU}((A, S, C), E_{\mathcal{P}_{\mathrm{Bel}}}, \mathbf{u}, \mathrm{Pl}_{\mathcal{P}_{\mathrm{Bel}}}) = \mathrm{CEU}(\mathcal{D}_*)$. Of course, $((A, S, C), E_{\mathcal{P}_{\mathrm{Bel}}}, \mathbf{u}, \mathrm{Pl}_{\mathcal{P}_{\mathrm{Bel}}}) \not\cong \mathcal{D}_*$; $\mathrm{Pl}_{\mathcal{P}_{\mathrm{Bel}}}$ and Bel are *not* the same, and they in fact represent related but different beliefs. (It is easy to show that sets are partially preordered by $\mathrm{Pl}_{\mathcal{P}_{\mathrm{Bel}}}$ but totally preordered by Bel.)

The key reason that GEU cannot represent nonuniform decision rules is because they do not respect the indistinguishability relations imposed by the utility function (and the plausibility measure). Recall that we require that *r{D) ~ "D* because we want a user of one decision rules to appear as if she were using another, without pretending that she has different tastes (and beliefs). So we want *r* to preserve the tastes (and beliefs) of its input.

There is a long-standing debate in the decision-theory literature as to whether preferences should be regarded as *ordinal* or *cardinal.* If they are ordinal, then all that matters is their order. If they are cardinal, then it should be meaningful to talk about the *differences* between preferences, that is, how much more a DM prefers one consequence to another. Similarly, if representations of likelihood are taken to be ordinal, then all that matters is whether one event is more likely than another. As we show below, if we require only that *r{V)* and *V* describe the same ordinal tastes (and beliefs), then we can in fact express almost all decision rules, including CEU, in terms of GEU.

Two utility functions $\mathbf{u}_1 : C \to U_1$ and $\mathbf{u}_2 : C \to U_2$ *represent the same ordinal tastes* if for all $c_1, c_2 \in C$, $\mathbf{u}_1(c_1) \precsim_{U_1} \mathbf{u}_1(c_2)$ iff $\mathbf{u}_2(c_1) \precsim_{U_2} \mathbf{u}_2(c_2)$. Similarly, two plausibility measures $\mathrm{Pl}_1 : 2^S \to P_1$ and $\mathrm{Pl}_2 : 2^S \to P_2$ *represent the same ordinal beliefs* iff for all $X, Y \subseteq S$, $\mathrm{Pl}_1(X) \preceq_{P_1} \mathrm{Pl}_1(Y)$ iff $\mathrm{Pl}_2(X) \preceq_{P_2} \mathrm{Pl}_2(Y)$. Finally, two decision problems $\mathcal{D}_1$ and $\mathcal{D}_2$ are *similar*, denoted $\mathcal{D}_1 \simeq \mathcal{D}_2$, iff they involve the same decision situations, their utility functions represent the same ordinal tastes, and their plausibility measures represent the same ordinal beliefs. Note that $\mathcal{D}_1 \cong \mathcal{D}_2$ implies $\mathcal{D}_1 \simeq \mathcal{D}_2$, but the converse is false in general. A decision rule transformation τ is an *ordinal $\mathcal{R}_1$ representation of $\mathcal{R}_2$* iff $\mathrm{dom}(\tau) = \mathrm{dom}(\mathcal{R}_2)$ and for all $\mathcal{D} \in \mathrm{dom}(\mathcal{R}_2)$, $\tau(\mathcal{D}) \simeq \mathcal{D}$ and $\mathcal{R}_1(\tau(\mathcal{D})) = \mathcal{R}_2(\mathcal{D})$.

We want to show next that almost all decision rules have an ordinal GEU representation. Doing so involves one more subtlety. Up to now, we have assumed that plausibility domains are partially ordered. This implies that two plausibility measures that represent the same ordinal beliefs necessarily induce the same indistinguishability relation (because of antisymmetry). Thus, in order to distinguish sets that have equivalent plausibilities when computing expected utility using $\otimes$ and $\oplus$, we need to allow plausibility domains to be partially *preordered.* So, for this result, we assume that $\precsim_P$ is a reflexive and transitive relation that is not necessarily antisymmetric (i.e., we could have that $p_1 \precsim_P p_2$ and $p_2 \precsim_P p_1$ but $p_1 \neq p_2$).

Theorem 3,6: *A decision rule R has an ordinal GEU representation iff R weakly respects utility.*

Proof: There are two cases, plausibilistic and nonplausibilistic. They are almost identical; we just do the plausibilistic case here. (Also, the "only if" direction is quite similar to the one in the oroof of Theorem 3.4. so we omit it here.')

Suppose that $\mathcal{R}$ is a plausibilistic decision rule that weakly respects utility. Fix a plausibilistic decision problem $\mathcal{D} = ((A, S, C), E_1, \mathbf{u}_1, \mathrm{Pl}_1) \in \mathrm{dom}(\mathcal{R})$. Let U_1 and P_1 be the utility domain and plausibility domain of E_1, respectively. Let $E_2 = (U_2, P_2, V, \oplus, \otimes)$ be defined as follows (note that $\times$ denotes Cartesian product in this proof):

- $U_2 = (U_1 \times C, \precsim_{U_2})$, where
  $$(u_1, c_1) \precsim_{U_2} (u_2, c_2) \text{ iff } u_1 \precsim_{U_1} u_2,$$
- $P_2 = (P_1 \times 2^S, \precsim_{P_2})$, where
  $$(p_1, X_1) \precsim_{P_2} (p_2, X_2) \text{ iff } p_1 \precsim_{P_1} p_2.$$
  (Note that $\precsim_{P_2}$ is a partial preorder, although it is not a partial order.)
- $V = (2^{S \times U_2}, \precsim_V)$, where $x \precsim_V y$ iff $x = y$ or
  1. $x = S \times \{(u_1, c_1)\}$, $y = S \times \{(u_2, c_2)\}$, and $(u_1, c_1) \precsim_{U_2} (u_2, c_2)$, or
  2. $x = \{(s, (\mathbf{u}_1(a(s)), a(s))) \mid s \in S\}$, $y = \{(s, (\mathbf{u}_1(b(s)), b(s))) \mid s \in S\}$, and $a \precsim_{\mathcal{R}(\mathcal{D})} b$, for some $a, b \in A$,
- $(p, X) \otimes (u, c) = X \times \{(u, c)\}$
- $x \oplus y = x \cup y$ for all $x, y \in V$.

Note that $(\perp_{P_1}, \emptyset) \precsim_{P_2} (p, X) \precsim_{P_2} (\top_{P_1}, S)$, so we have $\perp_{P_2} = (\perp_{P_1}, \emptyset)$ and $\top_{P_2} = (\top_{P_1}, S)$; thus, P_2 is a plausibility domain. Since $\mathcal{R}$ weakly respects utility, 1 and 2 are consistent with one another. We identify $(u, c) \in U_2$ with $S \times \{(u, c)\}$ in V; with this identification, $\top \otimes (u, c) = (u, c)$ for all $(u, c) \in U_2$ and, given 1 in the definition of $\precsim_V$, $(U_2, \precsim_{U_2})$ is a substructure of $(V, \precsim_V)$. Furthermore, $\oplus$ is clearly associative and commutative, so E_2 is indeed an expectation domain.

Now we need to define a utility function and a plausibility measure. Let $\mathbf{u}_2(c) = (\mathbf{u}_1(c), c)$ for all $c \in C$ and let $\mathrm{Pl}_2(X) = (\mathrm{Pl}_1(X), X)$ for all $X \subseteq S$. Note that

$$\mathrm{Pl}_2(X) \precsim_{P_2} \mathrm{Pl}_2(Y) \text{ iff } \mathrm{Pl}_1(X) \precsim_{P_1} \mathrm{Pl}_1(Y). \quad (3.6)$$

Thus Pl_2 is a plausibility measure, since Pl_1 is a plausibility measure. Also,

$$\mathbf{u}_2(c) \precsim_{U_2} \mathbf{u}_2(d) \text{ iff } \mathbf{u}_1(c) \precsim_{U_1} \mathbf{u}_1(d). \quad (3.7)$$

Let $\tau(\mathcal{D}) = ((A, S, C), E_2, \mathbf{u}_2, \mathrm{Pl}_2)$. Note that, by (3.6) and (3.7), $\tau(\mathcal{D}) \simeq \mathcal{D}$; furthermore, it is easy to check that $\mathbf{E}_{\mathrm{Pl}_2, E_2}((\mathbf{u}_2)_a) = \{(s, (\mathbf{u}_1(a(s)), a(s))) \mid s \in S\}$; it easily follows that $\mathrm{GEU}(\tau(\mathcal{D})) = \mathcal{R}(\mathcal{D})$. Thus τ is an ordinal GEU representation of $\mathcal{R}$. ∎

Theorem 3.6 shows that GEU can emulate essentially all decision rules. Thus, there is a sense in which GEU can be viewed as a universal decision rule. (We remark that although we have focused here on alternatives that are acts, in the sense of Savage, that is, functions from states to consequences, it is not hard to show—and we do in the full paper—that the same results hold if alternatives are taken to be horse lotteries, in the sense of Anscombe and Aumann [1963].)

References

[Anscombe and Aumann, 1963] F. J. Anscombe and R. J. Aumann. A definition of subjective probability. *Annals of Mathematical Statistics,* 34(1): 199-205, March 1963.

[Choquet, 1953] G. Choquet. Theory of capacities. *Annates de l'Institut Fourier (Grenoble),* 5:131-295, 1953.

[Chu and Halpern, 2003] F. C. Chu and J. Y. Halpern. Great Expectations. Part 1: On the customizability of generalized expected utility. In *Proc. Seventeenth International Joint Conference on Artificial Intelligence (IJCAI '01),* pages 35-40, 2003.

[Dempster, 1967] A. P. Dempster. Upper and lower probabilities induced by a multivalued mapping. *Annals of Mathematical Statistics,* 38:325-339, 1967.

[Fishburn, 1987] P. C. Fishburn. Reconsiderations in the foundations of decision under uncertainty. *The Economic Journal,* 97(388):825-841, December 1987.

[Friedman and Halpern, 1995] N. Friedman and J. Y. Halpern. Plausibility measures: a user's manual. In *Proc. Eleventh Conference on Uncertaintv in Artificial Intelligence (UAI '95),* pages 175-184. 1995.

[Giang and Shenoy, 2001] P. H. Giang and P. P. Shenoy. A comparison of axiomatic approaches to qualitative decision making using possibility theory. In *Proc. Seventeenth Conference on Uncertaintv in Artificial Intelligence (UAI 2001),* pages 162-170,2001.

[Gilboa and Schmeidler, 1989] 1. Gilboa and D. Schmeidler. Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics,* 18:141-153, 1989.

[Gul, 1991] F. Gul. A theory of disappointment aversion. *Econometrica,* 59(3):667-686, May 1991.

[Halpern, 2001] J. Y. Halpern. Conditional plausibility measures and Bayesian networks. *Journal of A.I. Research,* 14:359-389,2001.

[Lehmann, 2001] D. Lehmann. Expected qualitative utility maximization. *Games and Economic Behavior,* 35(1-2):54-79, April 2001.

[Quiggin, 1993] J. Quiggin. *Generalized Expected Utility Theory: The Rank-Dependent Expected Utility Model.* Kluwer Academic Publishers, Boston, Mass., 1993.

[Resnik, 1987] M. D. Resnik. *Choices; An Introduction to Decision Theory.* University of Minnesota Press, Minneapolis, 1987.

[Savage, 1954] L. J. Savage. *Foundations of Statistics.* John Wiley & Sons, New York, 1954.

[Schmeidler, 1989] D. Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57:571-587, 1989.

[Yaari, 1987] M. E. Yaari. The dual theory of choice under risk. *Econometrica,* 55(1):95-115, January 1987.