

Great Expectations.

Part I: On the Customizability of Generalized Expected Utility*

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

We propose a generalization of expected utility that we call *generalized EU* (GEU), where a decision maker's beliefs are represented by plausibility measures and the decision maker's tastes are represented by general (i.e., not necessarily real-valued) utility functions. We show that every agent, "rational" or not, can be modeled as a GEU maximizer. We then show that we can customize GEU by selectively imposing just the constraints we want. In particular, by show how of Savage's postulates corresponds to constraints on GEU.

1 Introduction

Many decision rules have been proposed in the literature. Perhaps the best-known approach is based on maximizing expected utility (EU). This approach comes in a number of variants; the two most famous are due to von Neumann and Morgenstern [1947] and Savage [1954]. They all follow the same pattern: they formalize the set of alternatives among which the decision maker (DM) must choose (typically as *acts*¹ or *lotteries*²). They then give a set of assumptions (often called *postulates* or *axioms*) such that the DM's preferences on the alternatives satisfy these assumptions iff the preferences have an EU representation, where an *EU representation* of a preference relation is basically a utility function (and a probability measure when acts are involved) such that the relation among the alternatives based on expected utility agrees with the preference relation. Moreover, they show that the representation is essentially unique. In other words, if the preferences of a DM satisfy the assumptions, then she is behaving *as if* she has quantified her tastes via a real-valued utility function (and her beliefs via a probability measure)

*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.

¹Formally, given a set S of states of the world and another set C of consequences, an *act* a is a function from S to C that, intuitively, associates with each state s the consequence of performing a in s .

²Formally, a *lottery* is a probability distribution over consequences; intuitively, the distribution quantifies how likely it is that each consequence occurs.

and she is relating the alternatives according to their expected utility. The assumptions are typically regarded as criteria for rational behavior, so these results also suggest that if a DM's beliefs are actually described by a probability measure and her tastes are described by a utility function, then she should relate the alternatives according to their expected utility (if she wishes to appear rational).

Despite the appeal of EU maximization, it is well known that people do not follow its tenets in general [Resnik, 1987]. As a result, a host of extensions of EU have been proposed that accommodate some of the more systematic violations (see, for example, [Gul, 1991; Gilboa and Schmeidler, 1989; Giang and Shenoy, 2001; Quiggin, 1993; Schmeidler, 1989; Yaari, 1987]). Again, the typical approach in the decision theory literature has been to prove representation theorems, showing that, given a suggested decision rule $\mathcal{R}$, there exists a collection of assumptions such that a preference relation satisfies the assumptions iff it has an $\mathcal{R}$ representation, where an $\mathcal{R}$ representation of a preference relation is essentially a choice of tastes (and beliefs) such that, given these as inputs, $\mathcal{R}$ relates the alternatives the way the preference relation does.

Given this plethora of rules, it would be useful to have a general framework in which to study decision making. The framework should also let us understand the relationship between various decision rules. We provide such a framework in this paper.

The basic idea of our approach is to generalize the notion of expected utility so that it applies in as general a context as possible. To this end, we introduce *expectation domains*, which are structures consisting of

- three (component) domains: a *plausibility domain*, a *utility domain*, and a *valuation domain*,
- two binary operators $\oplus$ and $\otimes$, which are the analogues of $+$ and $\times$ over the reals, and
- a reflexive binary relation $\preceq$ (which generalizes $\leq$).

Intuitively, $\otimes$ combines plausibility values and utility values much the same way that $\times$ combines probability and (real) utility, while $\oplus$ combines the products to form the (generalized) expected utility.

We have three domains because we do not want to require that DMs be able to add or multiply plausibility values or utility values, since these could be qualitative (e.g., plausibility

values could be {"unlikely", "likely", "very likely", ...} and utility values could be {"bad", "good", "better", ...}. In general, we do not assume that $\preceq$ is an order (or even a pre-order), since we would like to be able to represent as many preference relations and decision rules as possible.

Once we have an expectation domain, DMs can express their tastes and beliefs using components of the expectation domain. More specifically, the DMs express their beliefs using a *plausibility measure* [Friedman and Halpern, 1995], whose range is the plausibility domain of the expectation domain (plausibility measures generalize probability measures and a host of other representations of uncertainty, such as sets of probability measures, Choquet capacities, possibility measures, ranking functions, etc.) and they express their tastes using a utility function whose range is the utility domain of the expectation domain. In an expectation domain, it is possible to define a generalization of expected utility, which we call *generalized EU* (GEU). The GEU of an act is basically the sum (in the sense of $\oplus$) of products (in the sense of $\otimes$) of plausibility values and utility values that generalizes the standard definition of (probabilistic) expected utility over the reals in the obvious way.

We start by proving an analogue of Savage's result with respect to the decision rule (Maximizing) GEU.³ We show that every preference relations on acts has a GEU representation (even those that do not satisfy any of Savage's postulates), where a *GEU representation* of a preference relation basically consists of an expectation domain E , plausibility measure P , and utility function u , such that the way acts are related according to their GEU agrees with the preference relation (Theorem 3.1). In other words, no matter what the DM's preference relation on acts is, she behaves as if she has quantified her beliefs by a plausibility measure and her tastes via a utility function, and is relating the acts according to their (generalized) expected utility as defined by the $\oplus$ and $\otimes$ of some expectation domain. That is, we can model *any* DM using GEU, whether or not the DM satisfies any rationality assumptions. An important difference between our result and that of Savage is that he was constructing *EU* representations, which consists of a *real-valued* utility function u and a *probability* measure Pr (and the expectation domain is fixed, so $\oplus$, $\otimes$, and $\preceq$ are just $+$, $\times$, and $\leq$, respectively).

Given that GEU can represent *all* preference relations, it might be argued that GEU is too general—it offers no guidelines as to how to make decisions. We view this as a feature, not a bug, since our goal is to provide a general framework in which to express and study decision rules, instead of proposing yet another decision rule. Thus the absence of "guidelines" is in fact an absence of *limitations*: we do not want to exclude any possibilities at the outset, even preference relations that are not transitive or are incomplete. Starting from

³Many decision rules involve optimizing (i.e., maximizing or minimizing) some value function on the acts. Sometimes it is explicitly mentioned whether the function is to be maximized or minimized (e.g., "Minimax Regret" says explicitly to "minimize the maximum regret") while other times only the function name is mentioned and it is implicitly understood what is meant (e.g., "EU" means "maximize EU"). In this paper we will use "Maximizing GEU" and "GEU" interchangeably.

a framework in which we can represent all preference relations, we can then consider what preference relations have "special" representations, in the sense that the expectation domain, plausibility measure, and utility function in the representation satisfy some (joint) properties. This allows us to show how properties of expectation domains correspond to properties of preference relations. We can then "customize" GEU by placing just the constraints we want. We illustrate this by showing how each of Savage's postulates corresponds in a precise sense to an axiom on GEU.

Intuitively, a decision rule maps tastes (and beliefs) to preference relations on acts. Given two decision rules $\mathcal{R}_1$ and $\mathcal{R}_2$, an $\mathcal{R}_1$ *representation* of $\mathcal{R}_2$ is basically a function τ that maps inputs of $\mathcal{R}_2$ to inputs of $\mathcal{R}_1$ that represent the same tastes and beliefs, with the property that $\mathcal{R}_1(\tau(x)) = \mathcal{R}_2(x)$. Thus, τ models, in a precise sense, a user of $\mathcal{R}_2$ as a user of $\mathcal{R}_1$, since τ preserves tastes (and beliefs). In a companion paper [Chu and Halpern, 2003] we show that (almost) every decision rule has a GEU representation.

Although there has been a great deal of work on decision rules, there has been relatively little work on finding general frameworks for representing decision rules. In particular, there has been no attempt to find a decision rule that can represent all preference relations. There has been work in the fuzzy logic community on finding general notions of integration (which essentially amounts to finding notions of expectation) using generalized notions of $\oplus$ and $\otimes$; see, for example, [Benvenuti and Mesiar, 2000]. However, the expectation domain used in this work is (a subset of) the reals; arbitrary expectation domains are not considered. Luce [1990; 2000] also considers general addition-like operations applied to utilities, but his goal is to model joint receipts (which are typically modeled as commodity bundles in economics) as a binary operation, rather than to represent decision rules.

The rest of this paper is organized as follows. We cover some basic definitions in Section 2: plausibility domains, utility domains, expectation domains, decision problems, and GEU. We show that every preference relation on acts has a GEU representation in Section 3. In Section 4, we show that each of Savage's postulates corresponds to an axiom on GEU. We conclude in Section 5. Proofs of theorems stated are available at <http://www.cs.cornell.edu/home/halpern>.

2 Preliminaries

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. In particular, we do not want to force the DMs to linearly preorder all consequences and all events (i.e., subsets of the set of states). 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 $\perp_P$ and $\top_P$, such that for all $x \in P$, $\perp_P \preceq_P x \preceq_P \top_P$. Given a set S , a function $P1 : 2^S \rightarrow P$ is a *plausibility measure* iff

- PI1. $\text{Pl}(\emptyset) = \perp$,
 PI2. $\text{Pl}(S) = \top$, and
 PI3. if $X \subseteq Y$ then $\text{Pl}(X) \preceq \text{Pl}(Y)$.

Clearly plausibility measures are generalizations of probability measures. As pointed out in [Friedman and Halpern, 1995], plausibility measures generalize a host of other representations of uncertainty as well. Note that while the probability of any two sets must be comparable (since $\mathbb{R}$ is totally ordered), the plausibility of two sets may be incomparable.

We also want to represent the tastes of DMs using something more general than $\mathbb{R}$, so we allow the range of utility functions to be utility domains, where a *utility domain* is a set U endowed with a reflexive binary relation $\preceq_U$. Intuitively, elements of U represent the strength of likes and dislikes of the DM while elements of P represent the strength of her beliefs. Note that we do not require the DM's preference to be transitive (although we can certainly add this requirement). Experimental evidence shows that DMs' preferences occasionally do seem to violate transitivity.

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 $\times$).⁴ More formally, an *expectation domain* is a tuple $E = (U, P, V, \otimes, \oplus)$, where $(U, \preceq_U)$ is a utility domain, $(P, \preceq_P)$ is a plausibility domain, $(V, \preceq_V)$ is a valuation domain (where $\preceq_V$ is a reflexive binary relation), $\otimes : P \times U \rightarrow V$, and $\oplus : V \times V \rightarrow V$. (As usual, we omit subscripts when they are clear.) We have four requirements on expectation domains:

- E1. $(x \oplus y) \oplus z = x \oplus (y \oplus z)$;
 E2. $x \oplus y = y \oplus x$;
 E3. $\top \otimes x = x$;
 E4. $(U, \preceq_U)$ is a substructure of $(V, \preceq_V)$.

E1 and E2 say that $\oplus$ is associative and commutative. E3 says that $\top$ is the left-identity of $\otimes$ and E4 ensures that the expectation domain respects the relation on utility values.

Note that we do not require that $\oplus$ be monotonic; that is, we do not require that for all $x, y, z \in V$,

$$\text{if } x \preceq_V y \text{ then } x \oplus z \preceq_V y \oplus z. \quad (2.1)$$

We say that E is *monotonic* iff (2.1) holds. It turns out that monotonicity does not really make a difference by itself; see the comments after the proof of Theorem 3.1.

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 $\mathcal{A} = (A, S, C)$, 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).

⁴Sometimes we use $\times$ to denote Cartesian product; the context will always make it clear whether this is the case.

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 (*plausibilistic*) *decision problem* is a tuple $\mathcal{D} = (A, E, \mathbf{u}, \text{Pl})$, where

- $A = (A, S, C)$ is a decision situation,
- $E = (U, P, V, \otimes, \oplus)$ is an expectation domain,
- $\mathbf{u} : C \rightarrow U$ is a utility function, and
- $\text{Pl} : 2^S \rightarrow P$ is a plausibility measure.

We say that $\mathcal{D}$ is *monotonic* iff E is monotonic.

2.3 Expected Utility

Let $\mathcal{D} = ((A, S, C), E, \mathbf{u}, \text{Pl})$ be a plausibilistic decision problem. Each $a \in A$ induces a *utility random variable* $\mathbf{u}_a : S \rightarrow U$ as follows: $\mathbf{u}_a(s) = \mathbf{u}(a(s))$. In the standard setting (where utilities are real-valued 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:

$$\mathbf{E}_{\text{Pr}}(\mathbf{u}_a) = \sum_{x \in \text{ran}(\mathbf{u}_a)} \text{Pr}(\mathbf{u}_a^{-1}(x)) \times x. \quad (2.2)$$

We can generalize (2.2) 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

$$\mathbf{E}_{\text{Pl}, E}(\mathbf{u}_a) = \bigoplus_{x \in \text{ran}(\mathbf{u}_a)} \text{Pl}(\mathbf{u}_a^{-1}(x)) \otimes x. \quad (2.3)$$

We call (2.3) the *generalized EU* (GEU) of act a . Clearly (2.2) is a special case of (2.3).

In the probabilistic case, if all singleton sets are measurable with respect to Pr (i.e., in the domain of Pr), then

$$\mathbf{E}_{\text{Pr}}(\mathbf{u}_a) = \sum_{s \in S} \text{Pr}(s) \times \mathbf{u}_a(s). \quad (2.4)$$

The plausibilistic analogue of (2.4) is not necessarily equivalent to (2.3). A decision problem $((A, S, C), E, \mathbf{u}, \text{Pl})$ is *additive* iff

$$\text{Pl}(X \cup Y) \otimes \mathbf{u}(c) = \text{Pl}(X) \otimes \mathbf{u}(c) \oplus \text{Pl}(Y) \otimes \mathbf{u}(c)$$

for all $c \in C$ and nonempty $X, Y \subseteq S$ such that $X \cap Y = \emptyset$. Note that the notion of additivity we defined is a joint property of several components of a decision problem (i.e., $\oplus$, $\otimes$, $\mathbf{u}$, and Pl) instead of being a property of Pl alone. Additivity is exactly the requirement needed to make the analogue of (2.4) equivalent to (2.3). While decision problems involving probability are additive, those involving representations of uncertainty such as Dempster-Shafer belief functions are not, in general.

3 Representing Arbitrary Preference Relations

In this section, we show that every preference relation on acts has a GEU representation. GEU, like all decision rules, is formally a function from decision problems to preference relations on acts. Thus a GEU representation of a preference relation $\succsim_A$ on the acts in $\mathcal{A} = (A, \dots)$ is a decision problem $\mathcal{D} = (\mathcal{A}, E, \mathbf{u}, \text{Pl})$, where $E = (U, P, V, \oplus, \otimes)$, such that $a_1 \succsim_A a_2$ iff $\mathbf{E}_{\text{Pl}, E}(\mathbf{u}_{a_1}) \succsim_V \mathbf{E}_{\text{Pl}, E}(\mathbf{u}_{a_2})$.

Theorem 3.1: Every $\succsim_A$ has a GEU representation.⁵

Proof: Fix some $\mathcal{A} = (A, S, C)$ and $\succsim_A$. We want to construct a decision problem $\mathcal{D} = (\mathcal{A}, E, \mathbf{u}, \text{Pl})$ such that $\text{GEU}(\mathcal{D}) = \succsim_A$.

The idea is to let each consequence be its own utility and each set be its own plausibility, and define $\otimes$ and $\oplus$ such that each act is its own expected utility. For each $c \in C$, let a_c denote the constant act with the property that $a_c(s) = c$ for all $s \in S$. Let $E = (U, P, V, \oplus, \otimes)$ be defined as follows:

1. $U = (C, \succsim_C)$, where $c \succsim_C d$ iff $c = d$ or $a_c, a_d \in A$ and $a_c \succsim_A a_d$. (Note that Savage assumes that A contains all simple acts; in particular, A contains all constant acts. We do not assume that here.)
2. $P = (2^S, \subseteq)$.
3. $V = (2^{S \times C}, \succsim_V)$, where $x \succsim_V y$ iff $x = y$ or $x, y \in A$ and $x \succsim_A y$. (Note that set-theoretically a function is a set of ordered pairs, so $A \subseteq 2^{S \times C}$.)
4. $x \oplus y = x \cup y$ for $x, y \in V$.
5. $X \otimes c = X \times \{c\}$ for $X \in 2^S (= P)$ and $c \in C (= U)$.

We can identify $c \in C$ with $S \times \{c\}$ in V ; with this identification, $(U, \succsim_U)$ is a substructure of $(V, \succsim_V)$ and $\top \otimes c = c$ for all $c \in U (= C)$, as required. Furthermore, $\oplus$ is clearly associative and commutative, so E is indeed an expectation domain. Let $\mathcal{D} = (\mathcal{A}, E, \mathbf{u}, \text{Pl})$, where $\mathbf{u}(c) = c$ and $\text{Pl}(X) = X$. Note that

$$\begin{aligned} \mathbf{E}_{\text{Pl}, E}(\mathbf{u}_a) &= \bigoplus_{x \in \text{ran}(\mathbf{u}_a)} \text{Pl}(\mathbf{u}_a^{-1}(x)) \otimes x \\ &= \bigoplus_{c \in \text{ran}(a)} \text{Pl}(a^{-1}(c)) \otimes c \\ &= \{(s, c) \mid a(s) = c\} \\ &= a. \end{aligned}$$

That is, each act is its own expected utility; by the definition of $\succsim_V$, it is clear that $a \succsim_A b$ iff $\mathbf{E}_{\text{Pl}, E}(\mathbf{u}_a) \succsim_V \mathbf{E}_{\text{Pl}, E}(\mathbf{u}_b)$. Thus $\text{GEU}(\mathcal{D}) = \succsim_A$, as desired. ■

Note that the representation constructed in Theorem 3.1 is in fact additive, since if $X \cap Y = \emptyset$, then

$$(X \cup Y) \times \{c\} = (X \times \{c\}) \cup (Y \times \{c\}).$$

The representation is not necessarily monotonic. However, as we show in the full paper, a simple modification to the definition of $\succsim_V$ makes the representation monotonic, so monotonicity by itself does not restrict the kind of preference relations that can be represented.

⁵Note that, unlike most representation theorems, there is no uniqueness condition. However, we can show that our representation is canonical in a certain sense.

Theorem 3.1 holds in large part because of the flexibility we have. Given a decision situation (A, S, C) and a preference relation $\succsim_A$ on A , we are able to construct an expectation domain and a relation $\succsim_V$ that is customized to capture the relation $\succsim_A$ on A . However, this depends (in part) on two features of our setup.

The first is that, following Savage [1954], we took acts to be functions from states to consequences. Suppose instead that we have a *consequence function* $c : A \times S \rightarrow C$ that takes an act a and a state s and gives the consequence of a in s . Of course, in this setting, two distinct acts a_1 and a_2 could induce the same function from states to consequences; that is, we might have $c(a_1, s) = c(a_2, s)$ for all $s \in S$. It is easy to see that if a_1 and a_2 induce the same function from states to consequences, then no matter what expectation domain, utility function, and plausibility measure we use, a_1 and a_2 will have the same expected utility. Thus, if $\succsim_A$ does not treat a_1 and a_2 the same way, then $\succsim_A$ has no GEU representation.

A second reason that we do not need consistency constraints on $\succsim_A$ is that we have placed no constraints on $\succsim_V$, and relatively few constraints on $\oplus$, $\otimes$, $\mathbf{u}$, and Pl . If, for example, we required $\succsim_V$ to be transitive, then we would also have to require that $\succsim_A$ be transitive. The lack of constraints on $\oplus$, $\otimes$, $\mathbf{u}$, and Pl is important because it gives us enough freedom to ensure that distinct acts have different expected utility. In the next section, we investigate what happens when we add more constraints.

4 Representing Savage's Postulates

Theorem 3.1 shows that GEU can represent any preference relation. We are typically interested in representing preference relations that satisfy certain constraints, or postulates. The goal of this section is to examine the effect of such constraints on the components that make up GEU. For definiteness, we focus on Savage's postulates. For ease of exposition, we restrict to additive decision problems in this section; recall that this restriction does not affect Theorem 3.1.

A set $\mathcal{P}_e$ of axioms about (i.e., constraints on) plausibilistic decision problems *represents* a set of postulates $\mathcal{P}_r$ about decision situations and preference relations iff

$$\mathcal{D} = (\mathcal{A}, \dots) \text{ satisfies } \mathcal{P}_e \text{ iff } (\mathcal{A}, \text{GEU}(\mathcal{D})) \text{ satisfies } \mathcal{P}_r.$$

Theorem 3.1 can be viewed as saying that the empty set of axioms represents the empty set of postulates.

Before we present Savage's postulates, we first introduce some notation that will make the exposition more succinct. Suppose that $f : X \rightarrow Y$, $g : X \rightarrow Y$, and $Z \subseteq X$. Let $\langle f, Z, g \rangle$ denote the function h such that $h(x) = f(x)$ for all $x \in Z$ and $h(x) = g(x)$ for all $x \in \overline{Z}$. For example, if $X = Y = \mathbb{R}$ and $Z = \{x \mid x < 0\}$, then $\langle -x, Z, x \rangle$ is the absolute value function. In the intended application, the functions in question will be acts (i.e., functions from the set of states S to consequences C). So $a = \langle a_1, X, a_2 \rangle$ is the act such that $a(s) = a_1(s)$ for all $s \in X$ and $a(s) = a_2(s)$ for all $s \in \overline{X}$. For brevity, we identify the consequence $c \in C$ with the constant act a_c such that $a_c(s) = c$ for all $s \in S$. So for $c_1, c_2 \in C$, $\langle c_1, X, c_2 \rangle$ is the act with the property that $a(s) = c_1$ for all $s \in X$ and $a(s) = c_2$ for all $s \in \overline{X}$. Recall

that $X_1, \dots, X_n$ is a *partition* of Y iff the X_i 's are nonempty and pairwise disjoint, and $\bigcup_i X_i = Y$.

Fix some decision situation (A, S, C) . Readers familiar with [Savage, 1954] will recall that Savage implicitly assumes that A consists of all possible functions from S to C , since the DM can be questioned about any pair of functions. Throughout this paper, we have assumed that A could be any nonempty subset of the set of all simple acts. It is possible to maintain that assumption here, though some of the postulates would fail, not because $\lesssim_A$ does not relate certain members of A , but because A does not contain certain pairs of acts. We could adapt the postulates by “relativizing” them (so acts not in A are not required to be related) as we do in the full paper; however, that involves changing the statements of Savage’s postulates somewhat, and the presentation becomes more complicated. To simplify the exposition here, we assume in this section only that A consists of *all* simple acts. Here are Savage’s first six postulates (stated under the assumption that A consists of all simple acts from S to C , so that, for example, if a and b are in A , and $X \subseteq S$, then $(a, X, b) \in A$):

- P1. For all $a_1, a_2, a_3 \in A$,
- (a) $a_1 \lesssim_A a_2$ or $a_2 \lesssim_A a_1$, and
 - (b) if $a_1 \lesssim_A a_2$ and $a_2 \lesssim_A a_3$, then $a_1 \lesssim_A a_3$.
- P2. For all $a_1, a_2, b_1, b_2 \in A$ and $X \subseteq S$,
- $\langle a_1, X, b_1 \rangle \lesssim_A \langle a_2, X, b_1 \rangle$ iff $\langle a_1, X, b_2 \rangle \lesssim_A \langle a_2, X, b_2 \rangle$.
- P3. For all $X \subseteq S$, if there exist some $a_1, a_2 \in A$ such that for all $b \in A$, $\langle a_1, X, b \rangle \prec_A \langle a_2, X, b \rangle$, then for all $c_1, c_2 \in C$, $c_1 \lesssim_A c_2$ iff
- for all $a \in A$, $\langle c_1, X, a \rangle \prec_A \langle c_2, X, a \rangle$, or
 - for all $a \in A$, $\langle c_1, X, a \rangle \sim_A \langle c_2, X, a \rangle$.
- P4. For all $X_1, X_2 \subseteq S$, $c_1, d_1, c_2, d_2 \in C$, if $d_1 \prec_A c_1$ and $d_2 \prec_A c_2$, then
- $\langle c_1, X_1, d_1 \rangle \lesssim_A \langle c_1, X_2, d_1 \rangle$ iff $\langle c_2, X_1, d_2 \rangle \lesssim_A \langle c_2, X_2, d_2 \rangle$.
- P5. There exist $c_1, c_2 \in C$ such that $c_1 \prec_A c_2$.
- P6. For all $a, b \in A$, $c \in C$, if $a \prec_A b$, then there exists a partition $Z_1, \dots, Z_n$ of S , such that for all Z_i , $\langle c, Z_i, a \rangle \prec_A b$ and $a \prec_A \langle c, Z_i, b \rangle$.

P1 is the standard necessary condition for representation by EU (and many of its generalizations), since the reals are linearly ordered; it basically says that $\lesssim_A$ is a total preorder. More specifically, P1(a) says that $\lesssim_A$ is total (from which reflexivity follows) and P1(b) says that it is transitive. P2 is the *sure-thing principle*, which says that the way two acts are related depends only on their differences (that is, the part on which they agree can be ignored, since that is the “sure thing”). Savage defines for each subset $X \subseteq S$ a conditional preference relation as follows: $a_1 \lesssim_A^X a_2$ iff

- for all $a \in A$, $\langle a_1, X, a \rangle \prec_A \langle a_2, X, a \rangle$, or
- for all $a \in A$, $\langle a_1, X, a \rangle \sim_A \langle a_2, X, a \rangle$.

Intuitively, $a_1 \lesssim_A^X a_2$ iff when X occurs, the DM would find a_2 at least as good as a_1 . In the presence of P1, P2 holds iff $\lesssim_A^X$ is a total preorder for all X . Using $\lesssim_A^X$, Savage defines what it means for X to be null: X is *null* iff $a_1 \lesssim_A^X a_2$ for all $a_1, a_2 \in A$. P3 says that if X is not null, then $\lesssim_A^X$ and $\lesssim_A$ agree on the consequences. Savage defines a relation $\lesssim_S$ on events as follows: $X \lesssim_S Y$ iff for all $c, d \in C$, if $d \prec_A c$ then $\langle c, X, d \rangle \lesssim_A \langle c, Y, d \rangle$. Intuitively, we expect the DM to prefer a binary act that is more likely (according to her beliefs) to yield the more desirable consequence. P4, in the presence of P1–P3, basically ensures that $\lesssim_S$ is a total preorder. P5 asserts that S is not null. P1–P5 by themselves do not allow the construction of a unique EU representation.⁶ For this we need P6, which roughly says that for all pairs of acts $a, b \in A$ and consequences $c \in C$, if $a \prec_A b$ then we can partition S into events such that the DM’s preference about a and b is unaffected if c were to happen in any element of the partition. Savage also has a seventh postulate, but it is relevant only for general (nonsimple) acts; since we consider only simple acts, we omit it here.

Since we focus on additive decision problems in this section, we can define a notion of conditional GEU, which greatly simplifies the presentation of the GEU axioms. Given a plausibilistic decision problem $\mathcal{D} = ((A, S, C), E, \mathbf{u}, \text{P1})$ and $\emptyset \neq Z \subseteq S$, define the GEU of act a restricted to Z as follows:

$$\mathbf{E}_{\text{Pl}, E}(\mathbf{u}_a \upharpoonright Z) = \bigoplus_{x \in \mathbf{u}_a(Z)} \text{Pl}(\mathbf{u}_a^{-1}(x) \cap Z) \otimes x,$$

Note that $\mathbf{E}_{\text{Pl}, E}(\mathbf{u}_a \upharpoonright S) = \mathbf{E}_{\text{Pl}, E}(\mathbf{u}_a)$.

Since we are restricting to additive decision problems, it is easy to check that, given a partition $X_1, \dots, X_n$ of $Y \subseteq S$, we have that

$$\mathbf{E}_{\text{Pl}, E}(\mathbf{u}_a \upharpoonright Y) = \mathbf{E}_{\text{Pl}, E}(\mathbf{u}_a \upharpoonright X_1) \oplus \dots \oplus \mathbf{E}_{\text{Pl}, E}(\mathbf{u}_a \upharpoonright X_n).$$

Also, it is easy to check that for all nonempty proper subsets X of S ,

$$\mathbf{E}_{\text{Pl}, E}(\mathbf{u}_{(a_1, X, a_2)}) = \mathbf{E}_{\text{Pl}, E}(\mathbf{u}_{a_1} \upharpoonright X) \oplus \mathbf{E}_{\text{Pl}, E}(\mathbf{u}_{a_2} \upharpoonright \overline{X}).$$

Let $\mathcal{E}_{\mathcal{D}}(X) = \{\mathbf{E}_{\text{Pl}, E}(\mathbf{u}_a \upharpoonright X) \mid a \in A\}$. (We omit the subscript $\mathcal{D}$ if it is clear from context.) Intuitively, $\mathcal{E}_{\mathcal{D}}(X)$ consists of all the expected utility values of acts in A restricted to X . To simplify the statement of one of the axioms, let $\langle\langle u, X, v \rangle\rangle = \text{Pl}(X) \otimes u \oplus \text{Pl}(\overline{X}) \otimes v$ if $\emptyset \neq X \neq S$, and let $\langle\langle u, \emptyset, v \rangle\rangle = v$ and $\langle\langle u, S, v \rangle\rangle = u$.

A1. For all $x, y, z \in \mathcal{E}(S)$,

- (a) $x \lesssim_V y$ or $y \lesssim_V x$, and
- (b) if $x \lesssim_V y$ and $y \lesssim_V z$, then $x \lesssim_V z$.

A2. For all $\emptyset \neq X \neq S$, $x_1, x_2 \in \mathcal{E}(X)$, $y_1, y_2 \in \mathcal{E}(\overline{X})$,

- $x_1 \oplus y_1 \lesssim_V x_2 \oplus y_1$ iff $x_1 \oplus y_2 \lesssim_V x_2 \oplus y_2$.

A3. For all $\emptyset \neq X \neq S$, if there exist $x_1, x_2 \in \mathcal{E}(X)$ such that for all $y \in \mathcal{E}(\overline{X})$ $x_1 \oplus y \prec_V x_2 \oplus y$, then for all $u_1, u_2 \in \text{ran}(\mathbf{u})$, $u_1 \lesssim_V u_2$ iff

⁶They do, however, ensure that $\lesssim_S$ is a *qualitative probability* (which is a binary relation on 2^S with certain properties) [Savage, 1954].

- for all $y \in \mathcal{E}(\overline{X})$,
 $\text{Pl}(X) \otimes u_1 \oplus y \prec_V \text{Pl}(X) \otimes u_2 \oplus y$, or
- for all $y \in \mathcal{E}(\overline{X})$,
 $\text{Pl}(X) \otimes u_1 \oplus y \sim_V \text{Pl}(X) \otimes u_2 \oplus y$.

A4. For all $X_1, X_2 \subseteq S$, $u_1, v_1, u_2, v_2 \in \text{ran}(\mathbf{u})$,
 if $v_1 \prec_V u_1$ and $v_2 \prec_V u_2$, then

- $\langle\langle u_1, X_1, v_1 \rangle\rangle \precsim_V \langle\langle u_1, X_2, v_1 \rangle\rangle$ iff
 $\langle\langle u_2, X_1, v_2 \rangle\rangle \precsim_V \langle\langle u_2, X_2, v_2 \rangle\rangle$

A5. For some $u_1, u_2 \in \text{ran}(\mathbf{u})$, $u_1 \prec_V u_2$.

A6. For all $x, y \in \mathcal{E}(S)$, $u \in \text{ran}(\mathbf{u})$, if $x \prec_V y$, then
 for all $a, b \in A$, $c \in C$, such that $\mathbf{E}_{\text{Pl}, E}(\mathbf{u}_a) = x$,
 $\mathbf{E}_{\text{Pl}, E}(\mathbf{u}_b) = y$, and $\mathbf{u}(c) = u$, there exists a partition
 $Z_1, \dots, Z_n$ of S , such that x can be expressed as
 $x_1 \oplus \dots \oplus x_n$ and y can be expressed as $y_1 \oplus \dots \oplus y_n$,
 where $x_i = \mathbf{E}_{\text{Pl}, E}(\mathbf{u}_a \upharpoonright Z_i)$ and $y_i = \mathbf{E}_{\text{Pl}, E}(\mathbf{u}_b \upharpoonright Z_i)$ for
 $1 \leq i \leq n$, and for all $1 \leq k \leq n$,

- $\text{Pl}(Z_k) \otimes u \oplus \bigoplus_{i \neq k} x_i \prec_V y$, and
- $x \prec_V \text{Pl}(Z_k) \otimes u \oplus \bigoplus_{i \neq k} y_i$.

A1 says that the expected utility values are linearly preordered; more specifically, A1(a) says that they are totally preordered and A1(b) says that the relation is transitive. Note that A1 does not say that the whole valuation domain is linearly preordered: that would be a sufficient but not a necessary condition for GEU($\mathcal{D}$) to satisfy P1. Since we want necessary and sufficient conditions for our representation results, some axioms apply only to expected utility values rather than to arbitrary elements of the valuation domain. As the following theorem shows, every subset of $\{A1(a), A1(b), \dots, A6\}$ represents the corresponding subset of $\{P1(a), P1(b), \dots, P6\}$.

Theorem 4.1: For all $I \subseteq \{1(a), 1(b), \dots, 6\}$, $\{A_i \mid i \in I\}$ represents $\{P_i \mid i \in I\}$.

Theorem 4.1 is a strong representation result. For example, if we are interested in capturing all of Savage's postulates but the requirement that $\precsim_A$ is totally ordered, and instead are willing to allow it to be partially ordered (a situation explored by Lehmann [1996]), we simply need to drop axiom A1(b). Although we have focused here on Savage's postulates, it is straightforward to represent many of the other standard postulates considered in the decision theory literature in much the same way.

5 Conclusion

We have introduced GEU, a notion of generalized EU and shown that GEU can (a) represent all preference relations on acts and (b) be customized to capture any subset of Savage's postulates. In [Chu and Halpern, 2003], we show that GEU can be viewed as a universal decision rule in an even stronger sense: almost every decision rule has a GEU representation. Thus, the framework of expectation structures together with GEU provides a useful level of abstraction in which to study the general problem of decision making and rules for decision making.

References

- [Benvenuti and Mesiar, 2000] P. Benvenuti and R. Mesiar. Integrals with respect to a general fuzzy measure. In M. Grabisch, T. Murofushi, and M. Sugeno, editors, *Fuzzy Measures and Applications—Theory and Applications*, pages 205–232. Physica Verlag, Heidelberg, 2000.
- [Chu and Halpern, 2003] F. Chu and J. Y. Halpern. Great expectations. Part II: Generalized expected utility as a universal decision rule. In *IJCAI '03*, 2003.
- [Fishburn, 1987] P. C. Fishburn. Reconsiderations in the foundations of decision under uncertainty. *The Economic Journal*, 97(388):825–841, 1987.
- [Friedman and Halpern, 1995] N. Friedman and J. Y. Halpern. Plausibility measures: a user's guide. In *Proc. 11th Conf. Uncertainty in AI*, 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. 17th Conf. Uncertainty in AI*, pages 162–170, 2001.
- [Gilboa and Schmeidler, 1989] I. 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, 1991.
- [Lehmann, 1996] D. Lehmann. Generalized qualitative probability; Savage revisited. In *Proc. 12th Conf. Uncertainty in AI*, pages 318–388, 1996.
- [Luce, 1990] R. D. Luce. Rational versus plausible accounting equivalences in preference judgments. *Psychological Science*, 1:225–234, 1990.
- [Luce, 2000] R. D. Luce. *Utility of Gains and Losses: Measurement-Theoretical and Experimental Approaches*. Lawrence Erlbaum, London, 2000.
- [Neumann and Morgenstern, 1947] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 2nd edition, 1947.
- [Quiggin, 1993] J. Quiggin. *Generalized Expected Utility Theory: The Rank-Dependent Expected Utility Model*. Kluwer, 1993.
- [Resnik, 1987] M. D. Resnik. *Choices: An Introduction to Decision Theory*. University of Minnesota Press, 1987.
- [Savage, 1954] L. J. Savage. *Foundations of Statistics*. Wiley, 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, 1987.