

Great Expectations.

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

Francis C. Chu 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 [Chu and Halpern 2003] is a universal decision rule, in the sense that it can represent essentially all other decision rules. This approach gives us a general technique for designing new decision rules as well as providing a framework for comparing decision rules to each other.

1 Introduction

A great deal of effort has been devoted to studying decision making. A standard formalization describes 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; Kahneman and Tversky 1979; Luce 2000; Quiggin 1993; Schmeidler 1989; Tversky and Kahneman 1992; 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. A preliminary version of this paper appeared in *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI 2003)*, 2003.

In [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 \mathcal{R}_1 and \mathcal{R}_2 , an \mathcal{R}_1 *representation* of \mathcal{R}_2 is a function τ that maps inputs of \mathcal{R}_2 to inputs of \mathcal{R}_1 that contain the same representation of tastes (and beliefs) such 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). 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], that have no GEU representations.¹ This is because τ is not allowed to modify the representation of the tastes (and beliefs). We then define a notion of *ordinal representation*, in which τ 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.

We would like to emphasize again that it is important to distinguish the main result of [Chu and Halpern 2003], which shows that every *preference relation* has a GEU representation, from the results of this paper, which show that many *decision rules* have GEU representations and almost all decision rules have ordinal GEU representations. Representing a decision rule is not the same as representing a preference relation. Formally, a decision rule \mathcal{R} represents a preference relation \succsim if there exists some input x to \mathcal{R} (where x represents the tastes and perhaps beliefs of the DM) such that $\mathcal{R}(x) = \succsim$. On the other hand, \mathcal{R}_1 represents \mathcal{R}_2 if, roughly speaking, there is a *function* τ (rather than some input of \mathcal{R}_1) such that for *all* possible inputs x of \mathcal{R}_2 , $\mathcal{R}_1(\tau(x)) = \mathcal{R}_2(x)$. That is, $\tau \circ \mathcal{R}_1$ and \mathcal{R}_2 behave essentially the same way as functions on the domain of \mathcal{R}_2 . (Note that we can consider τ a *reduction* of \mathcal{R}_2 to \mathcal{R}_1 . Thus we can define and study hierarchies of decision rules, much the same way we can define and study hierarchies of languages and problems in the theory of computation. This topic, however, is beyond the scope of the paper.)

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 that he calls *Expected Qualitative Utility Maximization*, which allows utilities to be nonstandard real numbers 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.

The rest of this paper is organized as follows. We cover some basic definitions in Section 2: ex-

¹The CEU decision rule is the appropriate one to use if belief is represented by a Dempster-Shafer belief function; see Section 2.4 for more discussion.

pectation domains, decision problems, GEU, and decision rules (some of this material is taken from [Chu and Halpern 2003]). We show that most decision rules have (ordinal) GEU representations in Section 3, using Savage’s act framework. In Section 4, we show how these results can be applied to the lottery framework originally introduced by von Neumann and Morgenstern [1947] and the *horse lotteries* of Anscombe and Aumann [1963]. We conclude in Section 5 with some discussion of these results. Proofs are deferred to the appendix.

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 Halpern 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 \perp_P and \top_P , such that $\perp_P \preceq_P x \preceq_P \top_P$ for all $x \in P$. (We often omit the subscript P in \perp_P and \top_P when it is clear from context.) A function $\text{Pl} : 2^S \rightarrow P$ is a *plausibility measure* iff

$$\text{Pl1. } \text{Pl}(\emptyset) = \perp,$$

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

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

As pointed out in [Friedman and Halpern 1995], plausibility measures do not only probability, but also a host of other representations of uncertainty as well. 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.

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).² More formally, an *expectation domain* is a tuple $E = (U, P, V, \oplus, \otimes)$, 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), $\oplus : V \times V \rightarrow V$ and $\otimes : P \times U \rightarrow V$. We have four requirements on expectation domains:

$$E1. (x \oplus y) \oplus z = x \oplus (y \oplus z);$$

$$E2. x \oplus y = y \oplus x;$$

$$E3. \top \otimes x = x;$$

$$E4. (U, \preceq_U) \text{ is a substructure of } (V, \preceq_V).$$

$E1$ 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.

The *standard expectation domain*, which we denote \mathbb{E} , is $(\mathbb{R}, [0, 1], \mathbb{R}, +, \times)$, where the ordering on each domain is the standard order on the reals.

²We sometimes use \times to denote Cartesian product; the context will always make it clear whether this is the case.

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).

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 $(\mathcal{A}, U, \mathbf{u})$, where

- $\mathcal{A} = (A, S, C)$ is a decision situation,
- U is a utility domain, and
- $\mathbf{u} : C \rightarrow U$ is a utility function.

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

- $\mathcal{A} = (A, S, C)$ 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 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, and these components certainly can be ignored if they are not needed (see below).

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

2.3 Expected Utility

Let \mathcal{D} be a decision problem with S as the set of states, U as the utility domain, and \mathbf{u} as the utility function. Each act a of \mathcal{D} induces a *utility random variable* $\mathbf{u}_a : S \rightarrow U$ as follows: $\mathbf{u}_a(s) = \mathbf{u}(a(s))$. If in addition \mathcal{D} is plausibilistic with P as the plausibility domain and Pl as the plausibility measure, then each a also induces a *utility lottery* $\ell_a^{\text{Pl}, \mathbf{u}} : \text{ran}(\mathbf{u}_a) \rightarrow P$ as follows: $\ell_a^{\text{Pl}, \mathbf{u}}(u) = \text{Pl}(\mathbf{u}_a^{-1}(u))$. Intuitively, $\ell_a^{\text{Pl}, \mathbf{u}}(u)$ is the likelihood of getting utility u when performing act a . 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:

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

As we mentioned earlier, since acts are assumed to be simple, this sum is finite. We can generalize (1) to an arbitrary expectation domain $E = (U, P, V, \oplus, \otimes)$ by replacing $+$, \times , and Pr by \oplus , \otimes , and Pl , respectively. This gives us

$$\mathbf{E}_{\text{Pl}, E}(\mathbf{u}_a) = \bigoplus_{x \in \text{ran}(\mathbf{u}_a)} \text{Pl}(\mathbf{u}_a^{-1}(x)) \otimes x. \quad (2)$$

We call (2) the *generalized EU* (GEU) of act a . Clearly (1) is a special case of (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* \mathcal{R} is a function whose domain, denoted $\text{dom}(\mathcal{R})$, is a collection of (non)plausibilistic decision problems, and whose range, denoted $\text{ran}(\mathcal{R})$, is a collection of preference relations on acts.³ If $\mathcal{D} \in \text{dom}(\mathcal{R})$ and a_1 and a_2 are acts in \mathcal{D} , then we write

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

That is, taking \uparrow to denote “undefined”, if \mathcal{R} is a (non)plausibilistic decision rule and \mathcal{D} is a (non)plausibilistic decision problem, then $\mathcal{R}(\mathcal{D})$ is relation on the acts of \mathcal{D} if \mathcal{R} is defined on \mathcal{D} and $\mathcal{R}(\mathcal{D}) = \uparrow$ if \mathcal{R} is not defined on \mathcal{D} (e.g., if $\mathcal{R} = \text{Maximin}$ and the utility domain of \mathcal{D} is not linearly ordered). Here are a few examples of decision rules:

- GEU is a plausibilistic decision rule whose domain consists of all plausibilistic decision problems. Given a plausibilistic decision problem $\mathcal{D} = (\mathcal{A}, E, \mathbf{u}, \text{Pl})$, where $E = (U, P, V, \oplus, \otimes)$, we have

$$a_1 \succsim_{\text{GEU}(\mathcal{D})} a_2 \text{ iff } \mathbf{E}_{\text{Pl}, E}(\mathbf{u}_{a_1}) \succsim_V \mathbf{E}_{\text{Pl}, E}(\mathbf{u}_{a_2})$$

for all acts a_1, a_2 in \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).

³Readers familiar with set theory will note that the collection of all decision problems (plausibilistic or nonplausibilistic) is not a set, but a proper class. We can get around this problem by relativizing to sets, but this would complicate the presentation. For ease of exposition, we ignore the issue of proper classes in this paper.

- 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 picks 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 $\mathcal{D} = (\mathcal{A}, \mathbb{R}, \mathbf{u})$,

$$a_1 \succsim_{\text{Maximin}(\mathcal{D})} a_2 \text{ iff } \mathbf{w}_{\mathbf{u}}(a_1) \leq \mathbf{w}_{\mathbf{u}}(a_2).$$

Clearly the domain of Maximin can be extended so that it includes all nonplausibilistic decision problems where the range of the utility function is totally ordered.

- 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 $\mathcal{D} = ((A, S, C), \mathbb{R}, \mathbf{u})$, let $\bar{\mathbf{u}} : S \rightarrow U$ be defined as $\bar{\mathbf{u}}(s) = \sup_{a \in A} \mathbf{u}_a(s)$; that is, $\bar{\mathbf{u}}(s)$ is the least upper bound of the utilities in state s . The *regret* of a in state s , denoted $\mathbf{r}(a, s)$, is $\bar{\mathbf{u}}(s) - \mathbf{u}_a(s)$; note that no act can do better than a by more than $\mathbf{r}(a, s)$ in state s . Let $\bar{\mathbf{r}}(a) = \sup_{s \in S} \mathbf{r}(a, s)$. For example, suppose that $\bar{\mathbf{r}}(a) = 2$ and the DM picks a . Suppose that the DM then learns that the true state is s_0 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 than $\mathbf{u}_a(s_0)$. REG orders acts by their regret and thus takes the “best” act to be the one that minimizes $\bar{\mathbf{r}}(a)$. Intuitively, this rule tries to minimize the regret that a DM would feel if she discovered what the situation actually was: the “I wish I had chosen a_2 instead of a_1 ” feeling. Thus,

$$a_1 \succsim_{\text{REG}(\mathcal{D})} a_2 \text{ iff } \bar{\mathbf{r}}(a_1) \geq \bar{\mathbf{r}}(a_2).$$

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 \mathcal{P} 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 \mathcal{P}) 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 \mathcal{P} to $[0, 1]$, ordered pointwise; in other words, $p \preceq_P q$ iff $p(\text{Pr}) \leq q(\text{Pr})$ for all $\text{Pr} \in \mathcal{P}$. Thus, in this domain, \perp is the constant function 0 and \top is the constant function 1. For each $X \subseteq S$, let $f_X \in P$ be the function that evaluates each probability measure in \mathcal{P} at X ; that is, $f_X(\text{Pr}) = \text{Pr}(X)$ for all $\text{Pr} \in \mathcal{P}$. Let $\text{Pl}_{\mathcal{P}}(X) = f_X$; it is easy to verify that $\text{Pl}_{\mathcal{P}}$ is a plausibility measure. We view $\text{Pl}_{\mathcal{P}}$ as a representation of the set \mathcal{P} of probability measures; clearly \mathcal{P} can be recovered from $\text{Pl}_{\mathcal{P}}$. The domain of MMEU consists of all plausibilistic decision problems of the form $\mathcal{D} = ((A, S, C), (\mathbb{R}, [0, 1]^{\mathcal{P}}, V, \oplus, \otimes), \mathbf{u}, \text{Pl}_{\mathcal{P}})$, where \mathcal{P} is a set of probability measures on 2^S , and

$$a_1 \succsim_{\text{MMEU}(\mathcal{D})} a_2 \text{ iff } \inf_{\text{Pr} \in \mathcal{P}} \mathbf{E}_{\text{Pr}}(\mathbf{u}_{a_1}) \leq \inf_{\text{Pr} \in \mathcal{P}} \mathbf{E}_{\text{Pr}}(\mathbf{u}_{a_2}).$$

Note that this definition ignores \oplus , \otimes , and V .

- A *nonadditive probability* [Schmeidler 1989] is a function ν that associates each subset of a set S with a number between 0 and 1 such that $\nu(\emptyset) = 0$, $\nu(S) = 1$, and $\nu(X) \leq \nu(Y)$ if $X \subseteq Y$. (Roughly speaking, a nonadditive probability is just a plausibility measure whose range is $[0, 1]$, where $\perp = 0$ and $\top = 1$.) Schmeidler [1989] used a notion of expected utility for nonadditive probability that was defined by Choquet [1953]. (Choquet applied his notion of expectation to what he called *capacities*; nonadditive probabilities generalize capacities.) Given an act a , a real-valued utility function \mathbf{u} such that $\text{ran}(\mathbf{u}_a) = \{u_1, \dots, u_n\}$ and $u_1 < \dots < u_n$, and a nonadditive probability ν , define

$$\mathbf{E}_\nu(\mathbf{u}_a) = u_1 + \sum_{i=2}^n \nu(X_i) \times (u_i - u_{i-1}), \quad (3)$$

where $X_i = \mathbf{u}_a^{-1}(\{u_i, \dots, u_n\})$. It is easy to check (3) agrees with (1) if ν is a probability measure. The Choquet expected utility (CEU) rule has as its domain decision problems of the form $\mathcal{D} = (\mathcal{A}, \mathbb{E}, \mathbf{u}, \nu)$, and it orders acts as follows:

$$a_1 \succsim_{\text{CEU}(\mathcal{D})} a_2 \text{ iff } \mathbf{E}_\nu(\mathbf{u}_{a_1}) \leq \mathbf{E}_\nu(\mathbf{u}_{a_2}).$$

A special case of a nonadditive probability is a Dempster-Shafer *belief function* [Dempster 1967]. Belief functions also generalize probability. That is, every probability measure is a belief function, but not every belief function is a probability measure.⁴ Given a belief function Bel , it is well-known that there exists a set \mathcal{P}_{Bel} of probability measures such that for all $X \subseteq S$, $\text{Bel}(X) = \inf_{\text{Pr} \in \mathcal{P}_{\text{Bel}}} \text{Pr}(X)$ [Dempster 1967]. Moreover, if we use the CEU rule to compute expected belief, then it follows from results of Schmeidler [1986] that

$$\mathbf{E}_{\text{Bel}}(\mathbf{u}_a) = \inf_{\text{Pr} \in \mathcal{P}_{\text{Bel}}} \mathbf{E}_{\text{Pr}}(\mathbf{u}_a). \quad (4)$$

Let $\mathcal{D} = (\mathcal{A}, \mathbb{E}, \mathbf{u}, \text{Bel})$. It is immediate from (4) that if $\mathcal{D}_{\mathcal{P}_{\text{Bel}}}$ is the decision problem that results from \mathcal{D} by replacing Bel by $\text{Pl}_{\mathcal{P}_{\text{Bel}}}$ and replacing the plausibility domain $[0, 1]$ in the expectation domain by $[0, 1]^{\mathcal{P}_{\text{Bel}}}$, then $a_1 \succsim_{\text{CEU}(\mathcal{D})} a_2$ iff $a_1 \succsim_{\text{MMEU}(\mathcal{D}_{\mathcal{P}_{\text{Bel}}})} a_2$.⁵

3 Representing Decision Rules

Given a decision rule \mathcal{R} and a preference relation \succsim_A on the set of acts A , an \mathcal{R} *representation of* \succsim_A is basically a decision problem $\mathcal{D} \in \text{dom}(\mathcal{R})$ such that $\mathcal{R}(\mathcal{D}) = \succsim_A$ (and the set of acts in \mathcal{D} is A). In other words, an \mathcal{R} representation of \succsim_A makes \mathcal{R} relate acts in A the way \succsim_A relates them, so we can model a DM whose preference relation is \succsim_A as a user of \mathcal{R} . In [Chu and Halpern 2003] we prove the following:

Theorem 3.1 *Every preference relation \succsim_A has a GEU representation.*

We then go on to show how constraints on GEU can be used to capture various postulates on preference relations, such as Savage's postulates [Savage 1954].

⁴We assume that the reader is familiar with belief functions; see [Shafer 1976] for details. In any case, a knowledge of belief functions is not necessary for understanding the results of this paper.

⁵It follows from results of Schmeidler [1986] that a similar result holds, not just for belief functions, but for a larger set of nonadditive probability measures. Say that a probability measure Pr dominates a nonadditive probability ν on S if $\text{Pr}(X) \geq \nu(X)$ for all $X \subseteq S$. The result holds for all ν such that $\nu = \inf\{\text{Pr} \mid \text{Pr} \text{ dominates } \nu\}$.

In this paper, we go in a somewhat different direction. We start by extending 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 . We then investigate the extent to which GEU can represent arbitrary decision rules. 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 \succsim_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 $\text{dom}(\tau) = \text{dom}(\mathcal{R}_2)$ and for all $\mathcal{D} \in \text{dom}(\mathcal{R}_2)$,

- $\tau(\mathcal{D}) \cong \mathcal{D}$ 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 3.2 To see that Maximin has a GEU representation, let $E_{\max} = (\mathbb{R}, \{0, 1\}, \mathbb{R} \cup \{\infty\}, \min, \otimes)$, let Pl_{\max} be the plausibility measure such that $\text{Pl}_{\max}(X)$ is 0 if $X = \emptyset$ and 1 otherwise, and define $1 \otimes x = x$ and $0 \otimes x = \infty$. If $\mathcal{D} = (\mathcal{A}, \mathbb{R}, \mathbf{u})$, where $\mathcal{A} = (A, S, C)$, then it is easy to check that $\mathbf{E}_{\text{Pl}_{\max}, E_{\max}}(\mathbf{u}_a) = \mathbf{w}_{\mathbf{u}}(a)$. Take $\tau(\mathcal{D}) = (\mathcal{A}, E_{\max}, \mathbf{u}, \text{Pl}_{\max})$. Clearly $\tau(\mathcal{D}) \cong \mathcal{D}$: the decision situation and utility function have not changed. Moreover, it is immediate that $\text{GEU}(\tau(\mathcal{D})) = \text{Maximin}(\mathcal{D})$. ■

Example 3.3 To see that Minimax Regret (REG) has a GEU representation, for ease of exposition, we take $\text{dom}(\text{REG})$ to consist of standard decision problems $\mathcal{D} = ((A, S, C), \mathbb{R}, \mathbf{u})$ such that $M_{\mathcal{D}} = \sup_{s \in S} \bar{\mathbf{u}}(s) < \infty$. (If $M_{\mathcal{D}} = \infty$, given the restriction to simple acts, it is easy to show that all acts have infinite regret.) Let $E_{\text{reg}} = (\mathbb{R}, [0, 1], \mathbb{R} \cup \{\infty\}, \min, \otimes)$, where

$$x \otimes y = \begin{cases} y - \log(x) & \text{if } x > 0, \\ \infty & \text{if } x = 0. \end{cases}$$

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

$$\mathbf{E}_{\text{Pl}_{\mathcal{D}}, E_{\text{reg}}}(\mathbf{u}_a) = M_{\mathcal{D}} - \bar{\mathbf{r}}(a)$$

for all acts $a \in A$. Let $\tau(\mathcal{D}) = (\mathcal{A}, E_{\text{reg}}, \mathbf{u}, \text{Pl}_{\mathcal{D}})$. Clearly, $\tau(\mathcal{D}) \cong \mathcal{D}$, since the decision situation and utility function have not changed; furthermore, $\text{GEU}(\tau(\mathcal{D})) = \text{REG}(\mathcal{D})$, since higher expected utility corresponds to lower regret. ■

Example 3.4 To see that MMEU has a GEU representation, let $\mathcal{D} \in \text{dom}(\text{MMEU})$ such that $\mathcal{D} = (\mathcal{A}, (\mathbb{R}, [0, 1]^{\mathcal{P}}, \widehat{V}, \widehat{\oplus}, \widehat{\otimes}), \mathbf{u}, \text{Pl}_{\mathcal{P}})$. 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

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

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

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

Thus $\text{GEU}(\tau(\mathcal{D})) = \text{MMEU}(\mathcal{D})$; furthermore, it is clear that $\tau(\mathcal{D}) \cong \mathcal{D}$, since the decision situation, 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 \mathcal{R} respects utility if \mathcal{R} relates acts of constant utility according to the relation between utility values. Formally, a decision rule \mathcal{R} *respects utility* iff for all $\mathcal{D} \in \text{dom}(\mathcal{R})$ with A as the set of acts, S as the set of states, U as the utility domain, and \mathbf{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

$$a_1 \succsim_{\mathcal{R}(\mathcal{D})} a_2 \text{ iff } u_1 \succsim_U u_2. \quad (5)$$

We say that \mathcal{R} *weakly respects utility* iff (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, \succsim_U) is a substructure of (V, \succsim_V) . Thus if \mathcal{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 $\ell_{a_1}^{\text{Pl}, \mathbf{u}} = \ell_{a_2}^{\text{Pl}, \mathbf{u}}$,

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 $\mathcal{D} \in \text{dom}(\mathcal{R})$ and a_1, a_2, b_1, b_2 acts of \mathcal{D} such that $a_i \sim_{\mathcal{D}} b_i$,

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

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 \mathcal{R} is uniform, we can view \mathcal{R} as relating the acts indirectly by relating their descriptions.

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

Theorem 3.5 For all decision rules \mathcal{R} , \mathcal{R} has a GEU representation iff \mathcal{R} is uniform and \mathcal{R} respects utility.

Proof: See the appendix. ■

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

Example 3.6 Let $\mathcal{D}_* = ((A, S, C), \mathbb{E}, \mathbf{u}, \text{Bel})$, 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 $a_2(s_j) = 4 - j$, for $j = 1, 2, 3$; and
- Bel is the belief function such that $\text{Bel}(X) = 1$ if $\{s_1, s_2\} \subseteq X$ and $\text{Bel}(X) = 0$ otherwise.

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

$$\mathbf{E}_{\text{Bel}}(\mathbf{u}_{a_1}) = 1 + \text{Bel}(\{s_2, s_3\})(2 - 1) + \text{Bel}(\{s_3\})(3 - 2) = 1,$$

while

$$\mathbf{E}_{\text{Bel}}(\mathbf{u}_{a_2}) = 1 + \text{Bel}(\{s_1, s_2\})(2 - 1) + \text{Bel}(\{s_1\})(3 - 2) = 2.$$

It follows that CEU is not uniform, and so has no GEU representation. ■

How reasonable is the assumption of uniformity? That really depends on whether it is reasonable to identify two acts that are indistinguishable according to our definition. In the nonplausibilistic case, two acts are indistinguishable if, for all states s , the utility of their outcomes in state s are the same. If the utility of an act captures everything that is relevant about an act to the DM, then it does seem reasonable to say that two acts that are indistinguishable in this sense should be equally preferred by the DM. Arguably, if this is not the case, then the utility function is not capturing everything about the act that is important to the DM.

In the plausibilistic case, two acts a_1 and a_2 are indistinguishable if, roughly speaking, for each utility u , the likelihood of getting u according to a_1 is the same as the likelihood of getting u according to a_2 . However, it does *not* then in general follow that the likelihood of getting, say, either u_1 or u_2 according to a_1 is the same as the likelihood of getting either u_1 or u_2 according to a_2 . Decision rules whose input includes a plausibility measure where the likelihood of a set is not determined by the likelihood of its elements (note that belief functions are such plausibility measures) and whose behavior depends on the likelihood of obtaining one of a set of utilities, rather than just the likelihood of obtaining a single utility (CEU is such a rule) will not, in general, be uniform. Uniformity does not seem so compelling in this case though.

Can we say anything when uniformity does not hold? In fact, we can. To see why, first note that Example 3.4 shows that MMEU has a GEU representation. Moreover, as shown earlier, MMEU produces essentially the same order on acts as CEU restricted to belief functions. The fact that CEU has no GEU representation does not contradict Theorem 3.5. There is no decision problem \mathcal{D} such that $\mathcal{D} \cong \mathcal{D}_*$ (where \mathcal{D}_* is the decision problem in Example 3.6) and $\text{GEU}(\mathcal{D}) = \text{CEU}(\mathcal{D}_*)$. However, $\text{GEU}((A, S, C), E_{\mathcal{P}_{\text{Bel}}}, \mathbf{u}, \text{Pl}_{\mathcal{P}_{\text{Bel}}}) = \text{CEU}(\mathcal{D}_*)$. Of course, $((A, S, C), E_{\mathcal{P}_{\text{Bel}}}, \mathbf{u}, \text{Pl}_{\mathcal{P}_{\text{Bel}}}) \not\cong \mathcal{D}_*$; $\text{Pl}_{\mathcal{P}_{\text{Bel}}}$ 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 $\text{Pl}_{\mathcal{P}_{\text{Bel}}}$ 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 $\tau(\mathcal{D}) \cong \mathcal{D}$ 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 τ 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 $\tau(\mathcal{D})$ and \mathcal{D} 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 $\mathbf{u}_1 : C \rightarrow U_1$ and $\mathbf{u}_2 : C \rightarrow U_2$ *represent the same ordinal tastes* if for all $c_1, c_2 \in C$,

$$\mathbf{u}_1(c_1) \succsim_{U_1} \mathbf{u}_1(c_2) \text{ iff } \mathbf{u}_2(c_1) \succsim_{U_2} \mathbf{u}_2(c_2).$$

Similarly, two plausibility measures $\text{Pl}_1 : 2^S \rightarrow P_1$ and $\text{Pl}_2 : 2^S \rightarrow P_2$ *represent the same ordinal beliefs* iff for all $X, Y \subseteq S$,

$$\text{Pl}_1(X) \preceq_{P_1} \text{Pl}_1(Y) \text{ iff } \text{Pl}_2(X) \preceq_{P_2} \text{Pl}_2(Y).$$

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 $\text{dom}(\tau) = \text{dom}(\mathcal{R}_2)$ and for all $\mathcal{D} \in \text{dom}(\mathcal{R}_2)$,

- $\tau(\mathcal{D}) \simeq \mathcal{D}$ 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 computing expected utility using \oplus and \otimes , we need to allow plausibility domains to be partially *preordered*. So, for this result, we assume that \succsim_P is a reflexive and transitive relation that is not necessarily antisymmetric (i.e., we could have that $p_1 \succsim_P p_2$ and $p_2 \succsim_P p_1$ but $p_1 \neq p_2$).

Theorem 3.7 *A decision rule \mathcal{R} has an ordinal GEU representation iff \mathcal{R} weakly respects utility.*

Proof: See the appendix. ■

Theorem 3.7 shows that GEU can ordinally represent essentially all decision rules. Thus, there is a sense in which GEU can be viewed as a universal decision rule.

4 Related Frameworks

Note that so far we have worked exclusively in the *act framework* used by Savage [1954]. There are some other well-known frameworks in the decision-theory literature; perhaps the two best-known such frameworks are the *lottery framework* introduced by von Neumann and Morgenstern [1947], and Anscombe and Aumann’s [1963] *horse lotteries*, which can be viewed as a combination of the act and lottery frameworks. Since our goal is to provide a single framework for almost all of decision theory, in this section we briefly discuss how the act framework can model these, in much the same way as Turing machines can model other notions of computation. We begin with the lottery framework.

4.1 The Lottery Framework

As the name suggests, the alternatives in the lottery framework are lotteries, or probability distributions over consequences. Standard lotteries are functions of the form $\ell : C \rightarrow [0, 1]$ such that $\sum_{c \in C} \ell(c) = 1$. A standard lottery is *simple* iff $\{c \mid \ell(c) > 0\}$, which is typically called the *support* of ℓ and is denoted $\text{supp}(\ell)$, is finite. Note that the support of a standard lottery is nonempty.

In general, we want to allow assignments of plausibilities to sets of consequences. Given a set of consequences C and a plausibility domain P , a *lottery* is a plausibility measure $\ell : 2^Q \rightarrow P$, where Q is a nonempty subset of C . We denote Q as $\text{supp}(\ell)$. In the standard case, we take Q to consist of those consequences c such that $\ell(c) > 0$, so $\sum_{c \in \text{supp}(\ell)} \ell(c) = 1$. We say that ℓ is *degenerate* if $|\text{supp}(\ell)| = 1$, and we say that a lottery ℓ is *simple* iff $\text{supp}(\ell)$ is finite. Just as we focus on simple acts, we focus on simple lotteries (as did von Neumann and Morgenstern [1947]). Though lotteries are functions that assign plausibility values to consequences, we follow a common convention in the literature that lists plausibilities first (e.g., see [Sarin and Wakker 1997; Wakker and Zank 2002]). So $\{(p_1, c_1), \dots, (p_n, c_n)\}$ denotes the lottery ℓ such that $\text{supp}(\ell) = \{c_1, \dots, c_n\}$ and $\ell(c_i) = p_i$. (Note that this is the reverse of the usual notation for functions.)

Many notions we defined in the act framework have counterparts in the lottery framework. For example, the counterpart of a decision situation is a lottery decision situation. Formally, a *lottery decision situation* is a tuple $\mathcal{L} = (L, C, P)$, where

- C is a set of consequences,
- P is a plausibility domain, and
- L is a (nonempty) set of simple lotteries over C .

Note that a lottery decision situation does not contain any information about the tastes of the DM. A lottery decision problem is essentially a lottery decision situation together with information about the tastes of the DM. Formally, a *lottery decision problem* is a tuple $(\mathcal{L}, E, \mathbf{u})$, where

- $\mathcal{L} = (L, C, P)$ is a lottery decision situation,
- $E = (U, P, V, \oplus, \otimes)$ is an expectation domain, and
- $\mathbf{u} : C \rightarrow U$ is a utility function.

Note that the plausibility domain of the expectation domain is the same as the plausibility domain of the lottery decision situation.

A *standard lottery decision problem* is a lottery decision problem with the standard expectation domain; these are the ones that are studied extensively in the literature. Perhaps the best-known

lottery decision rule is von Neumann and Morgenstern’s expected utility rule: choosing the lottery that maximizes expected utility—that is, choosing the lottery ℓ that maximizes

$$\mathbf{E}_\ell(\mathbf{u}) = \sum_{c \in \text{supp}(\ell)} \ell(c) \times \mathbf{u}(c). \quad (6)$$

As in the act framework, we can generalize (6) to arbitrary expectation domains:

$$\mathbf{E}_{\ell,E}(\mathbf{u}) = \bigoplus_{c \in \text{supp}(\ell)} \ell(c) \otimes \mathbf{u}(c). \quad (7)$$

Some other well-known lottery decision rules include *disappointment aversion* [Gul 1991], *rank-dependent expected utility* [Quiggin 1993; Yaari 1987], and *cumulative prospect theory* [Tversky and Kahneman 1992; Wakker and Zank 2002]. The lottery framework has also been applied to non-probabilistic representations; for example, Giang and Shenoy [2001] give a representation theorem for lotteries based on possibility measures.

Our goal in this section is to show that the act framework can model the lottery framework. To facilitate this, we introduce one other notion in the act framework. A *plausibilistic decision situation* is a tuple $(\mathcal{A}, P, \text{Pl})$, where

- $\mathcal{A} = (A, S, C)$ is a decision situation,
- P is a plausibility domain, and
- $\text{Pl} : 2^S \rightarrow P$ is a plausibility measure.

Like a lottery decision situation, a plausibilistic decision situation describes the beliefs but not the tastes of the DM. The difference is, of course, that the belief of the DM is described by a single plausibility measure as opposed to a set of lotteries. Note that a plausibilistic decision problem is essentially a plausibilistic decision situation together with a utility function.

Given a plausibilistic decision situation $\mathcal{S} = ((A, S, C), P, \text{Pl})$, each $a \in A$ induces a lottery ℓ_a^{Pl} as follows: $\text{supp}(\ell_a^{\text{Pl}}) = \text{ran}(a)$ and $\ell_a^{\text{Pl}}(Y) = \text{Pl}(a^{-1}(Y))$ for $Y \subseteq \text{ran}(a)$. Note that if a is simple, then ℓ_a^{Pl} is also simple. We say that a plausibilistic decision situation \mathcal{S} *induces* the lottery decision situation $\mathcal{L}_\mathcal{S} = (\{\ell_a^{\text{Pl}} \mid a \in A\}, C, P)$.

This mapping from plausibilistic decision situations to lottery decision situations is clearly not 1-1. It is possible to have $\mathcal{S}_0 \neq \mathcal{S}_1$ but $\mathcal{L}_{\mathcal{S}_0} = \mathcal{L}_{\mathcal{S}_1}$, since different acts could induce the same lotteries (in fact, \mathcal{S}_0 and \mathcal{S}_1 may even involve different sets of states). However, as the following result shows, the mapping from plausibilistic decision situations to lottery decision situations is onto.

Proposition 4.1 *Every lottery decision situation $\mathcal{L} = (L, C, P)$ is induced by some plausibilistic decision situation $\mathcal{S}_\mathcal{L}$.*

Proof: See the appendix. ■

Corollary 4.2 *Every preference relation in the lottery framework can be modeled by a preference relation in the act framework.*

Proof: Let $\mathcal{S} = ((A, S, C), P, \text{Pl})$ be a plausibilistic decision situation and let $\mathcal{L} = (L, C, P)$ be the lottery decision situation it induces. Note that every preference relation \succsim_L on the lotteries in L induces a preference relation \succsim_A on the acts in A as follows:

$$a_1 \succsim_A a_2 \text{ iff } \ell_{a_1}^{\text{Pl}} \succsim_L \ell_{a_2}^{\text{Pl}}.$$

In other words, \succsim_A relates acts by the way \succsim_L relates the lotteries they induce. Since every lottery decision situation is induced by some plausibilistic decision situation (by Proposition 4.1), every preference relation in the lottery framework can be modeled in the act framework. ■

Note that an arbitrary preference relation \succsim_A on the acts in A does not correspond to a preference relation \succsim_L on the lotteries in L in general, since \succsim_A could treat acts that induce the same lottery differently. In order for \succsim_A to correspond to some \succsim_L , it must be *lottery-uniform*, in the sense that, for all acts a_1, a_2, b , if $\ell_{a_1}^{\text{Pl}} = \ell_{a_2}^{\text{Pl}}$, then

$$a_1 \succsim_A b \text{ iff } a_2 \succsim_A b \text{ and } b \succsim_A a_1 \text{ iff } b \succsim_A a_2.$$

It is not hard to see that lottery-uniform preference relations on acts are exactly those induced by preference relations on the lotteries. It is also not hard to see that not all preference relations on acts are lottery-uniform. So, some preferences that can be described by relating acts in a plausibilistic decision situation \mathcal{S} cannot be described by relating the lotteries in the lottery decision situation \mathcal{S} induces.

Turning now to decision problems and decision rules, we say that a plausibilistic decision problem $\mathcal{D}_A = (\mathcal{A}, E, \mathbf{u}, \text{Pl})$ induces the lottery decision problem $\mathcal{D}_L = (\mathcal{L}, E, \mathbf{u})$, where \mathcal{L} is the lottery decision situation induced by the the plausibilistic decision situation of \mathcal{D}_A , $(\mathcal{A}, P, \text{Pl})$. Since every plausibilistic decision problem induces a unique lottery decision problem, every lottery decision rule \mathcal{R}_L induces a plausibilistic decision rule \mathcal{R}_A as follows:

$$a_1 \succsim_{\mathcal{R}_A(\mathcal{D}_A)} a_2 \text{ iff } \ell_{a_1}^{\text{Pl}} \succsim_{\mathcal{R}_L(\mathcal{D}_L)} \ell_{a_2}^{\text{Pl}},$$

where \mathcal{D}_L is the lottery decision problem induced by \mathcal{D}_A . Basically, \mathcal{R}_A relates acts by relating the lotteries they induce using \mathcal{R}_L . The domain of \mathcal{R}_A is $\{\mathcal{D}_A \mid \mathcal{D}_A \text{ induces some } \mathcal{D}_L \in \text{dom}(\mathcal{R}_L)\}$. Thus every lottery decision rule can be modeled by a plausibilistic decision rule.

Using these observations, it is not hard to show that analogues of the results in previous sections also hold in the lottery framework. For example, it is easy to show that GEU when applied to lotteries yields a lottery decision rule that can represent all preference relation on lotteries and almost all lottery decision rules. More precisely, lottery GEU can represent all uniform lottery decision rules, where the notion of uniformity is completely analogous to the one presented in Section 3. In particular, it follows that lottery GEU can represent the well-known lottery decision rules mentioned earlier: disappointment aversion, rank-dependent expected utility, and cumulative prospect theory. They can also represent the rule considered by Giang and Shenoy based on possibility measures.

To summarize, all lottery decision rules can be modeled by plausibilistic decision rules. Thus it suffices, from a technical perspective, to focus exclusively on the act framework, as we have done in this paper, when considering the foundations of decision theory.

4.2 The Anscombe-Aumann Framework

Anscombe and Aumann [1963] define a framework that is essentially a combination of the act framework and the lottery framework: basically, it takes the consequences in the act framework and replaces them by lotteries, so acts (also known as *horse lotteries*) map states to lotteries (also known as *roulette lotteries*) The probabilities that the roulette lotteries assign to consequences are typically regarded as “objective” (in the sense that they are determined by the properties of the devices, such as fair coins or unloaded dice, used to generate them), while the probabilities (if any) associated with the sets of states are regarded, as in the act framework, as “subjective” (in the sense that these describe the beliefs of the DM).

We can formalize the AA framework in much the same way we formalized the act and lottery frameworks. As usual, we begin with decision situations. An *AA decision situation* is a tuple $\mathcal{H} = (H, S, \mathcal{L})$, where

- S is a set of states of the world,
- $\mathcal{L} = (L, C, P)$ is a lottery decision situation, and
- H is a nonempty set of horse lotteries (i.e., a nonempty subset of L^S).

A *nonplausibilistic AA decision problem* is a tuple $(\mathcal{H}, \widehat{E}, \mathbf{u})$, where

- $\mathcal{H} = (H, S, (L, C, \widehat{P}))$ is an AA decision situation,
- $\widehat{E} = (\widehat{U}, \widehat{P}, \widehat{V}, \widehat{\oplus}, \widehat{\otimes})$ is an expectation domain, and
- $\mathbf{u} : C \rightarrow \widehat{U}$ is a utility function.

Finally, a *plausibilistic AA decision problem* is a tuple $(\mathcal{H}, \widehat{E}, \mathbf{u}, E, \text{Pl})$, where

- $\mathcal{H} = (H, S, (L, C, \widehat{P}))$ is an AA decision situation,
- $\widehat{E} = (\widehat{U}, \widehat{P}, \widehat{V}, \widehat{\oplus}, \widehat{\otimes})$ is an expectation domain (for roulette lotteries),
- $\mathbf{u} : C \rightarrow \widehat{U}$ is a utility function,
- $E = (\widehat{V}, P, V, \oplus, \otimes)$ is an expectation domain (for horse lotteries), and
- $\text{Pl} : 2^S \rightarrow P$ is a plausibility measure.

We need two expectation domains, since in general the objective uncertainties and subjective uncertainties could be expressed in different languages. Note that the utility domain of E is the valuation domain of \widehat{E} , so the expected utility values with respect to the roulette lotteries are the utility values for E . While the formalization above is somewhat involved, in the standard setting, $\widehat{E} = \mathbb{E}$, and for the plausibilistic case, $E = \mathbb{E}$ as well.

In the standard setting, it is quite common to have the utility function map roulette lotteries, rather than just the (deterministic) consequences, to real numbers—that is, the domain of \mathbf{u} is L rather than C ; see, for example, [Anscombe and Aumann 1963; Gilboa and Schmeidler 1989; Schmeidler 1989]. This is because a utility function \mathbf{u} defined on C can easily be extended to L by taking $\mathbf{u}(\ell) = E_\ell(\mathbf{u})$. We can similarly extend \mathbf{u} to L in our framework, by taking $\mathbf{u}(\ell) = \mathbf{E}_{\ell, \widehat{E}}(\mathbf{u})$. Note that if ℓ is degenerate with $\text{supp}(\ell) = \{c\}$, then

$$\mathbf{u}(\ell) = \mathbf{E}_{\ell, \widehat{E}}(\mathbf{u}) = \top_{\widehat{P}} \widehat{\otimes} \mathbf{u}(c) = \mathbf{u}(c),$$

as one would expect.

Once we extend \mathbf{u} to L and treat lotteries as consequences, we can essentially view the AA framework as a special case of the act framework. As usual, a horse lottery h induces the random variable $\mathbf{u}_h : S \rightarrow \widehat{V}$ as follows: $\mathbf{u}_h(s) = \mathbf{u}(h(s))$. The expected utility of a horse lottery h is then

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

Thus, again, for the purpose of studying the foundations of decision theory, it suffices to focus on the act framework, since all decision rules in the AA framework can also be modeled by decision rules in the act framework.

5 Discussion

We have shown that (almost) all decision rules can be represented by GEU. So what does this result buy us? For one thing, decision rules are typically viewed as compact representations of how a DM makes decisions. Our results suggest a uniform way of representing decision rules that, in many cases of interest, will be compact. (How compact the representation is depends on how compactly we can describe \oplus , \otimes , and \succsim_V . While natural choices for these functions and relations typically do have a compact description, this is clearly not the case for all possible choices.) Moreover, our results provide a general technique for designing new decision rules, as well as providing a framework for comparing decision rules to each other. (As we observed in the introduction, we can in fact define a hierarchy on decision rules by treating representations as reductions.) This may be particularly relevant as we search for rules that are both adequate descriptively, in terms of describing what people actually do, and computationally tractable.

Acknowledgments

We thank Duncan Luce and Peter Wakker for useful comments on a draft of the paper.

A Proofs

Theorem 3.5 *For all decision rules \mathcal{R} , \mathcal{R} has a GEU representation iff \mathcal{R} is uniform and \mathcal{R} respects utility.*

Proof: We first show that if \mathcal{R} has a GEU representation, then it is uniform and respects utility. So, suppose that τ is a GEU representation of \mathcal{R} and let $\mathcal{D}_0 \in \text{dom}(\mathcal{R})$ 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 $\mathcal{D} = (\mathcal{A}, E, \mathbf{u}, \text{Pl}) \cong \mathcal{D}_0$, then $\mathbf{E}_{\text{Pl}, E}(\mathbf{u}_{a_i}) = \mathbf{E}_{\text{Pl}, E}(\mathbf{u}_{b_i})$. Thus for all plausibilistic \mathcal{D} , if $\mathcal{D} \cong \mathcal{D}_0$, then

$$a_1 \succsim_{\text{GEU}(\mathcal{D})} a_2 \text{ iff } b_1 \succsim_{\text{GEU}(\mathcal{D})} b_2.$$

Since τ is a GEU representation of \mathcal{R} , $\tau(\mathcal{D}_0) \cong \mathcal{D}_0$ and $\mathcal{R}(\mathcal{D}_0) = \text{GEU}(\tau(\mathcal{D}_0))$. It follows then that

$$a_1 \succsim_{\mathcal{R}(\mathcal{D}_0)} a_2 \text{ iff } b_1 \succsim_{\mathcal{R}(\mathcal{D}_0)} b_2;$$

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

$$a_1 \succsim_{\mathcal{R}(\mathcal{D}_0)} a_2 \text{ iff } a_1 \succsim_{\text{GEU}(\tau(\mathcal{D}_0))} a_2 \text{ iff } u_1 \succsim_U u_2,$$

where U is the utility domain of \mathcal{D}_0 , since τ is a GEU representation of \mathcal{R} . Thus \mathcal{R} respects utility.

We now show that, if \mathcal{R} is uniform and respects utility, then it has a GEU representation. We begin with the nonplausibilistic case.

Suppose that \mathcal{R} is a uniform nonplausibilistic decision rule that respects utility. Fix some decision problem $\mathcal{D} = ((A, S, C), U, \mathbf{u}) \in \text{dom}(\mathcal{R})$. Let $E = (U, P, V, \oplus, \otimes)$, where $P = (2^S, \subseteq)$, $V = 2^{S \times U}$, $x \oplus y = x \cup y$, and $X \otimes u = X \times \{u\}$. Now define \succsim_V as follows: $x \succsim_V y$ iff

1. $x = y$, or
2. $x = S \times \{u\}$ and $y = S \times \{v\}$ for some $u, v \in U$ such that $u \succsim_U v$, or

3. $x = \mathbf{u}_a$ and $y = \mathbf{u}_b$ for some $a, b \in A$ such that $a \succsim_{\mathcal{R}(\mathcal{D})} b$.

We need to check that \succsim_V is well defined. To see that clause 3 in the definition of \succsim_V does not introduce any inconsistencies by itself, we need to show that whenever we have $a_1, b_1, a_2, b_2 \in A$ such that $\mathbf{u}_{a_1} = \mathbf{u}_{a_2}$ and $\mathbf{u}_{b_1} = \mathbf{u}_{b_2}$, then $a_1 \succsim_{\mathcal{R}(\mathcal{D})} b_1$ iff $a_2 \succsim_{\mathcal{R}(\mathcal{D})} b_2$. Here is where we use the assumption that \mathcal{R} is uniform. Note that $\mathbf{u}_{a_1} = \mathbf{u}_{a_2}$ and $\mathbf{u}_{b_1} = \mathbf{u}_{b_2}$ implies that $a_1 \sim_{\mathcal{D}} a_2$ and $b_1 \sim_{\mathcal{D}} b_2$. Thus $a_1 \succsim_{\mathcal{R}(\mathcal{D})} b_1$ iff $a_2 \succsim_{\mathcal{R}(\mathcal{D})} b_2$, since \mathcal{R} is uniform. Note that clause 2 in the definition of \succsim_V essentially relates constant utility random variables; since \mathcal{R} respects utility, 2 and 3 are consistent with one another. Thus \succsim_V is well defined. We identify $u \in U$ with $S \times \{u\}$, so we have $\top \otimes u = u$, and it is clear that \oplus is associative and commutative. Given 2, it is easy to see that (U, \succsim_U) is a substructure of (V, \succsim_V) . Thus, E is an expectation domain.

Let $\text{Pl}(X) = X$ and $\tau(\mathcal{D}) = ((A, S, C), E, \mathbf{u}, \text{Pl})$. It is clear that $\tau(\mathcal{D}) \cong \mathcal{D}$, since the decision situation and utility function have not changed. Given the definitions of $\mathbf{E}_{\text{Pl}, E}(\mathbf{u}_a)$, E , Pl , and \mathbf{u} , we have

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

Given the definition of \succsim_V and the fact that $\mathbf{E}_{\text{Pl}, E}(\mathbf{u}_a) = \mathbf{u}_a$ for all $a \in A$, it is immediate that $\text{GEU}(\tau(\mathcal{D})) = \mathcal{R}(\mathcal{D})$. Thus τ is a GEU representation of \mathcal{R} .

The argument for the plausibilistic case is completely analogous, so we give a sketch here and leave the details to the reader. The key difference is that, instead of having $P = (2^S, \subseteq)$ and $\text{Pl}(X) = X$, the plausibility domain and plausibility measure are already given. So, instead of making $\mathbf{E}_{\text{Pl}, E}(\mathbf{u}_a) = \mathbf{u}_a$ (which is not possible in general, since we have to use the given plausibility measure), we make $\mathbf{E}_{\text{Pl}, E}(\mathbf{u}_a) = \ell_a^{\text{Pl}, \mathbf{u}}$; that is, $\mathbf{E}_{\text{Pl}, E}(\mathbf{u}_a)$ is the utility lottery induced by a instead of the utility random variable induced by a .

Suppose that \mathcal{R} is a uniform plausibilistic decision rule that respects utility. Fix some plausibilistic decision problem $\mathcal{D} = ((A, S, C), E_1, \mathbf{u}, \text{Pl}) \in \text{dom}(\mathcal{R})$. Let $E_2 = (U_1, P_1, V, \oplus, \otimes)$, where U_1 is the utility domain of E_1 , P_1 is the plausibility domain of E_1 , $V = 2^{P_1 \times U_1}$, $x \oplus y = x \cup y$, and $p \otimes u = \{(p, u)\}$. Define \succsim_V as follows: $x \succsim_V y$ iff

1. $x = y$, or
2. $x = \{(\top, u)\}$ and $y = \{(\top, v)\}$ for some $u, v \in U_1$ such that $u \succsim_{U_1} v$, or
3. $x = \ell_a^{\text{Pl}, \mathbf{u}}$ and $y = \ell_b^{\text{Pl}, \mathbf{u}}$ for some $a, b \in A$ such that $a \succsim_{\mathcal{R}(\mathcal{D})} b$.

Again, we need to check that \succsim_V is well defined. As in the nonplausibilistic case, it is easy to check that 3 does not introduce inconsistencies by itself, since \mathcal{R} is uniform. Also, since \mathcal{R} respects utility, 2 and 3 are consistent with one another. We identify $u \in U$ with $\{(\top, u)\}$, so $\top \otimes u = u$; given 2, it is easy to see that (U, \succsim_U) is a substructure of (V, \succsim_V) . Again, \oplus is associative and commutative. Thus E_2 is an expectation domain.

Let $\tau(\mathcal{D}) = (A, E_2, \mathbf{u}, \text{Pl})$. Obviously, $\tau(\mathcal{D}) \cong \mathcal{D}$, since the decision situation, utility function, and plausibility measure have not changed. It is easy to verify that $\mathbf{E}_{\text{Pl}, E_2}(\mathbf{u}_a) = \ell_a^{\text{Pl}, \mathbf{u}}$ for all $a \in A$. Thus it is immediate that $\text{GEU}(\tau(\mathcal{D})) = \mathcal{R}(\mathcal{D})$, given the definition of \succsim_V , so τ is a GEU representation of \mathcal{R} . ■

Theorem 3.7 *A decision rule \mathcal{R} has an ordinal GEU representation iff \mathcal{R} weakly respects utility.*

Proof: We first show that if \mathcal{R} has an ordinal GEU representation, then it is weakly respects utility. So, suppose that τ is an ordinal GEU representation of \mathcal{R} . Let $\mathcal{D}_1 \in \text{dom}(\mathcal{R})$ be arbitrary. Suppose that $a_{c_1} a_{c_2}$ are constant acts in \mathcal{D}_1 (where $a_{c_i}(s) = c_i$ for all states s). We need to show that

$$a_{c_1} \succsim_{\mathcal{R}(\mathcal{D}_1)} a_{c_2} \text{ iff } \mathbf{u}_1(c_1) \succsim_{U_1} \mathbf{u}_1(c_2),$$

where \mathbf{u}_1 is the utility function of \mathcal{D}_1 and U_1 is the utility domain of \mathcal{D}_1 . Let $\mathcal{D}_2 = \tau(\mathcal{D}_1)$; since τ is an ordinal GEU representation of \mathcal{R} , $\mathcal{D}_2 \simeq \mathcal{D}_1$ and $\text{GEU}(\mathcal{D}_2) = \mathcal{R}(\mathcal{D}_1)$. So

$$a_{c_1} \succsim_{\mathcal{R}(\mathcal{D}_1)} a_{c_2} \text{ iff } a_{c_1} \succsim_{\text{GEU}(\mathcal{D}_2)} a_{c_2} \text{ iff } \mathbf{u}_2(c_1) \succsim_{U_2} \mathbf{u}_2(c_2),$$

where \mathbf{u}_2 is the utility function of \mathcal{D}_2 and U_2 is the utility domain of \mathcal{D}_2 . Since $\mathcal{D}_2 \simeq \mathcal{D}_1$,

$$\mathbf{u}_2(c_1) \succsim_{U_2} \mathbf{u}_2(c_2) \text{ iff } \mathbf{u}_1(c_1) \succsim_{U_1} \mathbf{u}_1(c_2),$$

and we see that \mathcal{R} weakly respects utility.

Now we show that if \mathcal{R} weakly respects utility, then it has an ordinal GEU representation. As in Theorem 3.5, there are two cases, plausibilistic and nonplausibilistic. They are almost identical, so we do just the plausibilistic case here.

Suppose that \mathcal{R} is a plausibilistic decision rule that weakly respects utility. Fix a plausibilistic decision problem $\mathcal{D} = ((A, S, C), E_1, \mathbf{u}_1, \text{Pl}_1) \in \text{dom}(\mathcal{R})$. 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:

- $U_2 = (U_1 \times C, \succsim_{U_2})$, where $(u_1, c_1) \succsim_{U_2} (u_2, c_2)$ iff $u_1 \succsim_{U_1} u_2$.
- $P_2 = (P_1 \times 2^S, \succsim_{P_2})$, where $(p_1, X_1) \succsim_{P_2} (p_2, X_2)$ iff $p_1 \succsim_{P_1} p_2$. (Note that \succsim_{P_2} is a partial preorder, although it is not a partial order.)
- $V = (2^{S \times U_2}, \succsim_V)$, where $x \succsim_V y$ iff
 1. $x = y$, or
 2. $x = S \times \{(u_1, c_1)\}$, $y = S \times \{(u_2, c_2)\}$, and $(u_1, c_1) \succsim_{U_2} (u_2, c_2)$, or
 3. $x = \{(s, (\mathbf{u}_1(a(s)), a(s))) \mid s \in S\}$, $y = \{(s, (\mathbf{u}_1(b(s)), b(s))) \mid s \in S\}$, and $a \succsim_{\mathcal{R}(\mathcal{D})} b$, 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 $(\perp_{P_1}, \emptyset) \succsim_{P_2} (p, X) \succsim_{P_2} (\top_{P_1}, S)$, 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, clauses 2 and 3 in the definition of \succsim_V 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 $(u, c) \in U_2$ so it follows from clause 2 in the definition of \succsim_V that (U_2, \succsim_{U_2}) is a substructure of (V, \succsim_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 $\text{Pl}_2(X) = (\text{Pl}_1(X), X)$ for all $X \subseteq S$. Note that

$$\text{Pl}_2(X) \succsim_{P_2} \text{Pl}_2(Y) \text{ iff } \text{Pl}_1(X) \succsim_{P_1} \text{Pl}_1(Y). \quad (8)$$

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

$$\mathbf{u}_2(c) \succsim_{U_2} \mathbf{u}_2(d) \text{ iff } \mathbf{u}_1(c) \succsim_{U_1} \mathbf{u}_1(d). \quad (9)$$

Let $\tau(\mathcal{D}) = ((A, S, C), E_2, \mathbf{u}_2, \text{Pl}_2)$. Note that, by (8) and (9), $\tau(\mathcal{D}) \simeq \mathcal{D}$; furthermore, it is easy to check that $\mathbf{E}_{\text{Pl}_2, E_2}((\mathbf{u}_2)_a) = \{(s, (\mathbf{u}_1(a(s)), a(s))) \mid s \in S\}$; so $\text{GEU}(\tau(\mathcal{D})) = \mathcal{R}(\mathcal{D})$, given the definition of \succsim_V . Thus τ is an ordinal GEU representation of \mathcal{R} . ■

Proposition 4.1 *Every lottery decision situation $\mathcal{L} = (L, C, P)$ is induced by some plausibilistic decision situation $\mathcal{S}_{\mathcal{L}}$.*

Proof: We first prove the proposition for the standard case. Suppose that $\mathcal{L} = (L, C, [0, 1])$. Let $S = [0, 1)$. Suppose that $\ell \in L$ and $\text{supp}(\ell) = \{c_1^\ell, \dots, c_k^\ell\}$. Let a_ℓ be defined as follows: $a_\ell(s) = c_k^\ell$ for all $s \in S$ such that $\sum_{i=1}^{k-1} \ell(c_i^\ell) \leq s < \sum_{i=1}^k \ell(c_i^\ell)$. Let $\mathcal{S}_{\mathcal{L}} = ((A_L, S, C), [0, 1], \text{Pr})$, where

- Pr is the uniform distribution on S and
- $A_L = \{a_\ell \mid \ell \in L\}$.

It is easy to check that $\ell_{a_\ell}^{\text{Pr}} = \ell$, so $\mathcal{S}_{\mathcal{L}}$ induces \mathcal{L} .

The construction is more complicated for general plausibility domains, since we must make sure S is rich enough to allow us to use a single plausibility measure to induce all the lotteries. Given a lottery decision situation $\mathcal{L} = (L, C, P)$, let $S_L = \{f \mid f \in C^L \text{ and } f(\ell) \in \text{supp}(\ell)\}$. Intuitively, each state f assigns to each lottery ℓ some consequence in $\text{supp}(\ell)$. Let a_ℓ be defined by taking $a_\ell(f) = f(\ell)$. Now we need to specify a plausibility measure. The idea is to construct Pl so that $\text{Pl}(a_\ell^{-1}(X)) = \ell(X)$. Clearly this guarantees that $\ell_{a_\ell}^{\text{Pl}}(X) = \ell(X)$ for all $X \in 2^{\text{supp}(\ell)}$, so that a_ℓ induces ℓ .

To make the definition of Pl more concise, let $\varphi(\ell, Y)$ be the following statement: there exists some nonempty $X \subseteq \text{supp}(\ell)$ such that $a_\ell^{-1}(X) \subseteq Y$. Given $Y \subseteq S_L$, we define $\text{Pl}(Y)$ as follows:

1. If there does not exist $\ell \in L$ such that $\varphi(\ell, Y)$, let $\text{Pl}(Y) = \perp$.
2. If there exists a unique $\ell \in L$ such that $\varphi(\ell, Y)$, let $\text{Pl}(Y) = \ell(Z)$, where

$$Z = \bigcup \{X \mid X \subseteq \text{supp}(\ell) \text{ and } a_\ell^{-1}(X) \subseteq Y\}.$$

3. If there exist two distinct $\ell_1, \ell_2 \in L$ such that $\varphi(\ell_1, Y)$ and $\varphi(\ell_2, Y)$, let $\text{Pl}(Y) = \top$.

Note that for each $Y \subseteq S_L$, exactly one of the three cases applies, so Pl is well defined.

To see that Pl is a plausibility measure, note that clearly $\text{Pl}(S_L) = \top$ (since $L \neq \emptyset$) and $\text{Pl}(\emptyset) = \perp$. Now suppose that $Y_1 \subseteq Y_2$. We have three cases:

- Case 1 applies to Y_2 . Then it must apply to Y_1 as well, so $\text{Pl}(Y_1) = \perp \preceq \text{Pl}(Y_2)$.
- Case 2 applies to Y_2 ; let ℓ_2 be the unique lottery such that $\varphi(\ell_2, Y_2)$. Since $Y_1 \subseteq Y_2$, for all $\ell \in L$, $\varphi(\ell, Y_1)$ implies $\varphi(\ell, Y_2)$. Thus, if there is some $\ell \in L$ such that $\varphi(\ell, Y_1)$, it must be ℓ_2 . So either case 1 applies to Y_1 , then we are done as above, or $\varphi(\ell_2, Y_1)$. Since $Y_1 \subseteq Y_2$, if $a_{\ell_2}^{-1}(X) \subseteq Y_1$ then $a_{\ell_2}^{-1}(X) \subseteq Y_2$; thus $Z_1 \subseteq Z_2$, where

$$Z_i = \bigcup \{X \mid X \subseteq \text{supp}(\ell_2) \text{ and } a_{\ell_2}^{-1}(X) \subseteq Y_i\},$$

and so $\text{Pl}(Y_1) = \ell_2(Z_1) \preceq \ell_2(Z_2) = \text{Pl}(Y_2)$.

- Case 3 applies to Y_2 . Then $\text{Pl}(Y_1) \preceq \top = \text{Pl}(Y_2)$.

So Pl is a plausibility measure.

Now we want to show that $\text{Pl}(a_\ell^{-1}(X)) = \ell(X)$ for all $X \subseteq \text{supp}(\ell)$. Clearly this is true if $X = \emptyset$ or $X = \text{supp}(\ell)$. So suppose that X is a nonempty proper subset of $\text{supp}(\ell)$. Note that $\varphi(\ell, a_\ell^{-1}(X))$, so either case 2 or case 3 of the definition of Pl applies. Suppose that $\varphi(\ell_0, a_{\ell_0}^{-1}(X))$ for some $\ell_0 \in L$. Then there exists some nonempty $X_0 \subseteq \text{supp}(\ell_0)$ such that $a_{\ell_0}^{-1}(X_0) \subseteq a_{\ell_0}^{-1}(X)$. We want to show that $\ell_0 = \ell$, so that case 2 applies. Note that there exists some $c \in \text{supp}(\ell) - X$ and there exists some $c_0 \in X_0$ by assumption. Suppose that $\ell_0 \neq \ell$; then there exists some $f \in S_L$ such that $f(\ell) = c$ and $f(\ell_0) = c_0$ by construction. However, it is clear that $f \in a_{\ell_0}^{-1}(X_0)$ and $f \notin a_\ell^{-1}(X)$. Since $a_{\ell_0}^{-1}(X_0) \subseteq a_\ell^{-1}(X)$, no such f exists; it follows that $\ell_0 = \ell$ and so case 2 applies. Thus $\text{Pl}(a_\ell^{-1}(X)) = \ell(X)$ and so $\mathcal{S}_\mathcal{L}$ induces \mathcal{L} . ■

References

- Anscombe, F. J. and R. J. Aumann (1963). A definition of subjective probability. *Annals of Mathematical Statistics* 34(1), 199–205.
- Choquet, G. (1953). Theory of capacities. *Annales de l'Institut Fourier (Grenoble)* 5, 131–295.
- Chu, F. and J. Y. Halpern (2003). Great expectations. Part I: On the customizability of generalized expected utility. In *Proc. Eighteenth International Joint Conference on Artificial Intelligence (IJCAI '03)*, pp. 291–296.
- Dempster, A. P. (1967). Upper and lower probabilities induced by a multivalued mapping. *Annals of Mathematical Statistics* 38, 325–339.
- Fishburn, P. C. (1987). Reconsiderations in the foundations of decision under uncertainty. *The Economic Journal* 97(388), 825–841.
- Friedman, N. and J. Y. Halpern (1995). Plausibility measures: a user's guide. In *Proc. Eleventh Conference on Uncertainty in Artificial Intelligence (UAI '95)*, pp. 175–184.
- Giang, P. H. and P. P. Shenoy (2001). A comparison of axiomatic approaches to qualitative decision making using possibility theory. In *Proc. Seventeenth Conference on Uncertainty in Artificial Intelligence (UAI 2001)*, pp. 162–170.
- Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics* 18, 141–153.
- Gul, F. (1991). A theory of disappointment aversion. *Econometrica* 59(3), 667–686.
- Halpern, J. Y. (2001). Conditional plausibility measures and Bayesian networks. *Journal of A.I. Research* 14, 359–389.
- Kahneman, D. and A. Tversky (1979). Prospect theory: an analysis of decision under risk. *Econometrica* 47(2), 263–292.
- Lehmann, D. (2001). Expected qualitative utility maximization. *Games and Economic Behavior* 35(1–2), 54–79.
- Luce, R. D. (2000). *Utility of Gains and Losses: Measurement-Theoretical and Experimental Approaches*. London: Lawrence Erlbaum.
- Quiggin, J. (1993). *Generalized Expected Utility Theory: The Rank-Dependent Expected Utility Model*. Boston: Kluwer.

- Resnik, M. D. (1987). *Choices: An Introduction to Decision Theory*. Minneapolis: University of Minnesota Press.
- Sarin, R. and P. Wakker (1997). A single-stage approach to Anscombe and Aumann's expected utility. *The Review of Economic Studies* 64(3), 399–409.
- Savage, L. J. (1954). *Foundations of Statistics*. New York: Wiley.
- Schmeidler, D. (1986). Integral representation without additivity. *Proc. of the Amer. Math. Soc.* 97(2), 255–261.
- Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica* 57, 571–587.
- Shafer, G. (1976). *A Mathematical Theory of Evidence*. Princeton, N.J.: Princeton University Press.
- Tversky, A. and D. Kahneman (1992). Advances in prospect theory: cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5, 297–323.
- von Neumann, J. and O. Morgenstern (1947). *Theory of Games and Economic Behavior*. Princeton: Princeton University Press.
- Wakker, P. and H. Zank (2002). A simple preference foundation of cumulative prospect theory with power utility. *European Economic Review* 46(7), 1253–1271.
- Yaari, M. E. (1987). The dual theory of choice under risk. *Econometrica* 55(1), 95–115.