Ambiguity in the Context of Probabilistic Beliefs*

Klaus Nehring

University of California, Davis†

November 2001‡

ABSTRACT

When are a decision maker's preferences compatible with specified probabilistic beliefs? We propose two definitions of "compatibility" as possible answers to this question. The weaker one requires simply that the decision maker prefer to bet on an event whenever it is more likely. It is shown that under regularity conditions there exists a unique (incomplete) maximal comparative likelihood relation such that a given preference ordering is compatible with it. This relation yields a model-free notion of "revealed unambiguous beliefs"; it implies a definition of "revealed unambiguous events" that is strictly more demanding than Epstein-Zhang (2001)'s.

A stronger notion of compatibility called Tradeoff Consistency requires that the decision maker rank acts according to their expected utility whenever possible. Given "minimally complete" probabilistic beliefs, Tradeoff Consistency entails a complete determination of preferences over multi-valued acts by betting preferences and consequence utilities. Exluding all departures from expected utility maximization due to Allais-type probabilistic risk-aversion, it leads to a concept of "utility sophistication" dual to Machina-Schmeidler's "probabilistic sophistication". We also propose two "Ellsbergian" definitions of ambiguity aversion in terms of betting preferences and characterize their close link to the major notions of ambiguity aversion proposed in the literature.

*This version is still quite rough and incomplete; it is also clearly too long, and will most likely be split in future versions. Nonetheless, I would greatly appreciate any comments and suggestions.

†e-mail: kdnehring@ucdavis.edu

‡Some of the main ideas of this paper were presented in December 1996 at LOFT2 in Torino, Italy, and in July 2001 at RUD in Venice, Italy. I am grateful to Juergen Eichberger, Paolo Ghirardato and Clemens Puppe for helpful comments.

1. INTRODUCTION

In this paper, we shall study decision makers who have precise probabilistic beliefs about some events, while their beliefs about others may be imprecise or "ambiguous". Behaviorally, this ambiguity is characterized by violations of the sure-thing principle exemplified by the Ellsberg's (1961) celebrated (thought) experiments. While this and much other evidence provide compelling reasons to abandon the assumption that behavior can be globally explained in terms of precise probabilistic beliefs, they do not render the notion of probabilistic belief useless if it is applied "locally", that is: if applied to some events or event comparisons. Indeed, the very formulation of Ellsberg's original experiment suggests a comparison of events with probabilistic beliefs to those where beliefs are presumed to be ambiguous. In many situations in which ambiguity is plausible and interesting, the existence of local probabilistic beliefs is plausible as well. For example, while the expected return on equity (as a whole) is notoriously hard to pin down through statistical information, this is much less true for the volatility of the return. Thus, it seems much more plausible to assume precise probabilistic beliefs on the latter than on the former. Likewise, in the context of games under incomplete information, it may well make sense to ascribe to players precise beliefs about others' types along with ambiguity about their strategy choices.

To explore how precise and imprecise probabilistic beliefs are intertwined, and how they together determine an agent's preferences, our inquiry will take two directions that mirror each other. First, we shall ask when an agent's preferences are "compatible" with a specified set of probabilistic beliefs. We shall then ask, conversely, what probabilistic beliefs can meaningfully attributed to that agent on the basis of an observation of his preference ordering.

Two notions of compatibility of preferences with probabilistic beliefs will be proposed, a "minimalist" one and a "maximalist" one. Pertaining to preferences over bets only, the minimalist notion leaves maximal room for phenomena of "probabilistic risk-aversion"[1] as displayed in the famous Allais "paradox". By contrast, the maximalist notion excludes such phenomena systematically; it requires that preferences over multi-valued acts are based on comparisons of expected utility, whenever the agent's beliefs are sufficiently precise to allow this. If the DM has a sufficiently rich set of probabilistic beliefs, the stronger notion implies that the entire preference ordering is determined by preferences over bets and consequence utilities. Such preferences will be called "utility

[1] In the following, we shall use the name "probabilistic risk-averison" as a catch-all phrase for departures from expected utlity maximization with known probabilities.

sophisticated", in analogy to Machina-Schmeidler's (1992) notion of "probabilistic sophistication".

Imprecise Qualitative Probability

To lay the groundwork for the analysis, we shall describe a decision maker's (DM's) probabilistic beliefs as an incomplete "more-likely-than relation" over events $\trianglerighteq$ that can be represented in terms of a set of priors Π as follows:

$$A \trianglerighteq B \text{ if and only if } \pi(A) \geq \pi(B) \text{ for all } \pi \in \Pi. \tag{1}$$

The existence of this multi-prior representation ensures that the relation $\trianglerighteq$ fully incorporates the "logical syntax of probability", exemplified by entailments such as "if A is more likely than B, then not-B must be more likely than not-A". To obtain existence and in particular uniqueness of the representation, it is assumed that $\trianglerighteq$ is "minimally complete" in an appropriate sense. Minimally complete belief relations $\trianglerighteq$ with the multi-prior representation (1) will be referred to as "imprecise qualitative probabilities"; their characterization is proved and further developed in Nehring (2001).

Compatibility of Betting Preferences with Probabilistic Beliefs

A bet on event A is the act that engenders the better of the two consequences in the event A, and the worse consequence otherwise. A DM's "preferences over bets" consists of all preference comparisons of bets involving the same two consequences. A DM's betting preferences are *compatible* with his probabilistic information $\trianglerighteq$ if he prefers to bet on A to betting on B whenever A is more likely than B in terms of $\trianglerighteq$. Compatibility has bite through the consistency conditions defining an imprecise qualitative probability.

Throughout most of the paper, it will be assumed that the DM's preferences are compatible with a "minimally complete" belief relation $\trianglerighteq$ imprecise qualitative probability relation.[2] This structural assumption proves to be of great analytical and unifying power.[3] It is satisfied, for example, if preferences are compatible with probabilistic beliefs reflecting the existence of an independent, continuous random device as assumed in the Anscombe-Aumann approach.

[2]This relation is to be understood non-exhaustively, i.e. as capturing *some* of the probabilisitc beliefs that can be attributed to the decision maker but not necessarily all of them.

[3]It is discussed in section 5 and relinquished in section 7.

A Model-Free Definition of Revealed Unambiguous Beliefs

From a descriptive, third-person point of view it is of interest to reverse this question and ask what probabilistic beliefs can be meaningfully attributed to a DM's on the basis of observations of his preferences. To obtain a well-defined answer, one would like to be able to come up with a unique *maximal* imprecise qualitative probability such that the DM's preferences are compatible with it. This would allow a clear-cut answer to any question of the form: does the DM believe that A is more likely than B. It would then be meaningful to identify this maximal imprecise qualitative probability with the DM's "revealed unambiguous (probabilistic) beliefs". The first main result of the paper, Theorem 3, shows that a unique maximal imprecise qualitative probability exists under regularity conditions. The resulting notion of "revealed unambiguous beliefs" is model-free in that it makes no assumptions on probabilistic risk-attitudes, and only weak regularity assumptions on ambiguity attitudes. It gives rise to a definition of "revealed unambiguous events" as those for which the DM has a precise unambiguous belief; the definition turns out to be strictly more demanding than Epstein-Zhang's (2001) as it applies to betting preferences (cf. Proposition 1).[4]

The published literature has not addressed the issue of compatibility explicitly, as far as we know. Implicitly, however, it offers proposals for the special case of unconditional probabilistic beliefs through primitive[5] definitions of "unambiguous events" revealed by the preference relation. As explained in more detail at the end of the introduction, the compatibility requirements derived from the extant definitions fail to adequately capture the "syntax of probability", as in the case of Epstein-Zhang (2001) and Ghirardato-Marinacci (2001a)[6], or are of restricted applicability.

Utility Sophistication based on Tradeoff Consistency

While descriptively the case for phenomena of probabilistic risk-aversion is strong, it is much weaker normatively, even if violations of the sure-thing principle due to ambiguity are deemed normatively permissible. Moreover, for purposes of economic modelling, there is a clear interest in

[4]Epstein-Zhang (2001) build probabilistic sophistication over unambiguous acts into their definition of a "revealed unambiguous event"; we obtain this as an implication of a separate rationality axiom on preferences over multi-valued acts called "Revealed Stochastic Dominance"; see section 4.1.

[5]That is: definitions that are not based on a prior notion of compatibility.

[6]By contrast, the more recent work in progress Ghirardato-Maccheroni-Marinacci (2001c) proposes a definition of revealed unambiguous events that is effectively equivalent to that of Nehring (1999), and coincides with the one proposed here if preferences are "utility sophisticated". See section 7 for further discussion.

"zooming in" on ambiguity-based departures from SEU by excluding all others. Hence a stronger notion of compatibility with given probabilistic beliefs requires, intuitively speaking, that acts are compared in terms of expected utility "whenever possible". The stronger notion entails EU maximization over unambiguous acts, but goes substantially further. To flesh it out, assume that the DM's preferences over unambiguous acts have an EU representation with von Neumann-Morgenstern utility function u. Then the stronger requirement can be formulated as the condition that an act f be preferred to an act g whenever, given any "revealed admissible" prior $\pi \in \Pi$, the expected utility of f (with respect to π) exceeds that g. In section 4, this requirement is captured by an axiom of Tradeoff Consistency[7].

Under the maintained assumption that preferences are compatible with a minimally complete probabilistic belief relation $\trianglerighteq$, Tradeoff Consistency implies that the agent's preference relation over multi-valued acts is completely determined by preferences over bets and unambiguous acts. This is shown by the second main result of the paper, Theorem 4. In view of this determination, we shall call such preferences "utility sophisticated", intending this term to be dual to "probabilistically sophistication" in the sense of Machina-Schmeidler (1992): while the latter excludes ambiguity, but is largely unconstrained towards phenomena of probabilistic risk-aversion, the former excludes those systematically, but imposes only minimal constraints on the structure of ambiguity.

The best-known example of utility sophisticated preferences is the "minimum expected utility" (MEU) model due to Gilboa-Schmeidler (1989); in the MEU model, acts are evaluated according to $\min_{\pi \in \Pi'} E_\pi u \circ f$, for an appropriate convex set of priors Π'. Let ρ denote a representation of the agent's betting preferences that is normalized to be additive on unambiguous events. Then a utility sophisticated preference relation is MEU if and only if ρ is a lower probability, that is: if $\rho(A) = \min_{\pi \in \Pi'} \pi(A)$ for all events A. In section 6, we characterize betting preferences based on lower probabilities (and therefore of the MEU model under utility sophistication) using a novel notion of ambiguity aversion that is defined in terms of generalized Ellsberg-style experiments. Being a property of betting preferences, ambiguity aversion in our sense is clearly distinguished from phenomena of probabilistic risk-aversion. By contrast, as pointed out by Epstein (1999), the existing definitions of ambiguity aversion that give rise to the MEU model (essentially Schmeidler's (1989) original definition and versions thereof) fail to accommodate this distinction.

Adapting recent work by Ghirardato-Maccheroni-Marinacci (2001c), we also characterize the more

[7]Tradeoff Consistency is related to Wakker's (1989) axiom of "Noncontradictory Revealed Tradeoffs"; it is strong enough to entail EU maximization over unambiguous acts in the first place.

general class of preferences with a constant "degree of caution" γ (the "α-MEU model" in their terminology); this model is given by a representation of the form $\rho(A) = \gamma \min_{\pi \in \Pi^*} \pi(A) + (1 - \gamma) \max_{\pi \in \Pi^*} \pi(A)$. Here, utility sophistication entails an evaluation of acts according to $\gamma \min_{\pi \in \Pi^*} E_\pi u \circ f + (1 - \gamma) \max_{\pi \in \Pi^*} E_\pi u \circ f$.

By contrast, utility sophistication is incompatible with the widely used Choquet Expected Utility (CEU) model: CEU preferences are utility sophisticated only if the DM maximizes subjective expected utility. We interpret this initially perhaps somewhat disturbing state of affairs as saying that while the CEU model may be appropriate for subspaces on which the DM's beliefs have a rather specific structure, they are inappropriate if applied globally to preferences that are compatible with minimally complete probabilistic information; similar arguments suggesting limitations of the CEU model have in fact been made before in Klibanoff (2001a,b) and Nehring (1999).

Compatibility of Betting Preferences with Probabilistic Beliefs: An Example

The notion of compatibility of betting preferences with beliefs is non-trivial mainly due to the fact that precise beliefs over some events may entail imprecise restrictions on others, as illustrated by the following example.

Suppose that the decision maker knows the marginal distribution of each of two random variables to be uniform on the unit interval, but that she has little idea about their joint distribution. Consider her preferences over bets on events in the state space $\Theta = [0,1] \times [0,1]$; in the "bet on A" she receives 1 dollar if the event A occurs, and 0 otherwise. Her betting preferences are described by a set function $\rho : 2^\Theta \to [0,1]$ with the property that $\rho([a,b] \times [0,1]) = \rho([0,1] \times [a,b]) = b - a$, reflecting her knowledge of the marginal distributions; here $\rho(A)$ can be interpreted as the "probability equivalent of A", i.e. as the known probability to which the bet on A is equivalent in terms of preference. An elementary computation shows that, under any joint probability distribution π consistent with the given marginal distributions, the probability of the event $[0,c] \times [0,d]$ cannot be less than $c + d - 1$, and cannot exceed $c + d$. Our basic rationality postulate "Compatibility" requires that the bet on $[0,c] \times [0,d]$ is strictly to be preferred to the bet on any event T with known probability less than $c + d - 1$, and that it is strictly inferior to the bet on any event T with known probability exceeding $c + d$, i.e. that

$$c + d - 1 \leq \rho([0,c] \times [0,d]) \leq c + d, \tag{2}$$

whatever else she might believe, and whatever her ambiguity attitude is.[8]

As mentioned above, various definitions of "revealed unambiguous events" in terms of the DM's preferences can be found in the literature. Given such a definition of "revealed unambiguous events", one can define a preference relation to be compatible with a given set of (unconditional) probabilistic beliefs over events in some specified set Λ if the probability equivalents for all events in Λ are equal to the given probabilistic beliefs, and if all events in Λ are revealed unambiguous. In the example, Λ consists of all "marginal" events, i.e. all events of the form $A_1 \times [0,1]$ and $[0,1] \times A_2$.

Applying this notion of compatibility to the definitions of "revealed unambiguous events" proposed by Epstein-Zhang (2001) and Ghirardato-Marinacci (2001a), it is not difficult to construct betting preferences that are compatible with the given probabilistic information on marginals in terms of these proposals but nonetheless conflict with the compatibility restriction (2). This happens, for instance, with betting preferences described by the following set function ρ_*, with

$$\rho_*(A) = \sup\{\lambda(B) \mid B \times [0,1] \subseteq A \text{ or } [0,1] \times B \subseteq A\}, \text{ for any } A \subseteq \Theta,$$

where λ is the Lebesgue-measure on $[0,1]$.[9] It is not difficult to verify that this preference relation reveals all marginal events as unambiguous according to the Epstein-Zhang definition.[10] Yet ρ_* is far from satisfying the compatibility requirement (2), since $\rho_*([0,c] \times [0,d]) = 0$ whenever $c < 1$ and $d < 1$. Thus, while ρ_* fails to build in the restrictions on ambiguous events entailed by Compatibility, this deficiency is not recognized by the Epstein-Zhang definition. If the example is appropriately extended to render the alternative definition proposed by Ghirardato-Marinacci (2001a) applicable, that definition turns out to founder in the same way. The Ghirardato-Marinacci definition has two further limitations: first, it applies only to either uniformly ambiguity-averse or uniformly ambiguity-loving preferences (in a sense specified by them). Moreover, it assumes that all departures from SEU maximization are due to ambiguity, and thus does not allow for the type of "probabilistic risk-aversion" evidenced most famously in the Allais paradox.

A third definition of "revealed unambiguous beliefs" has been proposed in Nehring (1999). It was designed to induce the right notion of compatibility, and indeed performs without problems in this example; however, while applicable to arbitrary ambiguity attitudes, it is restrictive just like the

[8] In other words, even for the most ambiguity averse decision maker, the probability equivalent of an event could not be less than its lowest possible probability; and even for the most ambiguity loving decision maker, the probability equivalent of an event could not exceed its highest possible probability.

[9] ρ_* is in fact an "inner measure" of the kind considered in Epstein-Zhang (2001)

[10] This follows from the fact that, for any T of the form $[a,b] \times [0,1]$ and any A disjoint from T, $\rho_*(A \cup T) = \rho_*(A) + \rho_*(T)$.

Ghirardato-Marinacci definition by assuming that all departures from SEU maximization are due to ambiguity.

Finally, all published contributions are limited to unconditional probabilistic beliefs, but do not allow to incorporate conditional probabilistic information. Such types of information are fundamental, however. For example, standard statistical models typically consist of families of precise conditional distributions (frequently indexed by appropriate parameters); a natural interpretation of classical (frequentist) statistics in the tradition of Wald is to assume probabilistic ignorance (or at least imprecision) about the parameters themselves. Learning under ambiguity will thus frequently have to incorporate conditional probabilistic information.

A way of overcoming this limitation has first been proposed in the talk Nehring (1996), from which the present paper originated. Its fundamental idea of defining a "revealed unambiguous preference" relation as the maximal independent subrelation has recently been taken up and developed further by Ghirardato-Maccheroni-Marinacci (2001c). In particular, they use this relation to provide a characterization of preferences with a constant "degree of caution" mentioned above, and analyze the relation of preferences and revealed beliefs in a dynamic context.

Overview

Section 2 describes a DM's probabilistic beliefs in terms of a "comparative likelihood relation" on events, and states a multi-prior characterization result that is proved and further developed in Nehring (2001). Section 3 focuses on the relation between probabilistic beliefs and betting preferences. We first define a notion of compatibility of a DM's betting preferences with a specified comparative likelihood relation, and then specify conditions under which there exists a maximal such relation so that his preferences are compatible with it (Theorem 3). This relation is interpreted as representing the DM's "revealed unambiguous beliefs". A definition of "revealed unambiguous events" is derived from it and compared to that of Epstein-Zhang (2001). Section 4 then considers preferences over multi-valued acts; we first propose a "minimal" rationality axiom "Revealed Stochastic Dominance" that entails probabilistic sophistication of preferences over unambiguous acts. Strengthening Revealed Stochastic Dominance to Tradeoff Consistency, we then characterize the class of "utility sophisticated" preferences in which all departures from SEU are due to ambiguity (Theorem 4) . Sections 3 and 4 assume that the DM's preferences are compatible with a given minimally complete set of probabilistic beliefs. Section 5 asks whether such reference can be eliminated, and a "fully behavioral" account of revealed unambiguous beliefs can be given. We argue that the answer is positive if preferences are utility sophisticated, but negative otherwise. Section 6 presents two "Ellsbergian" definitions of ambiguity aversion; the weaker corresponds to Epstein's (1999) definition, while the stronger yields a characterization of the MEU model without confounding ambiguity aversion with probabilistic risk-aversion. Finally, in section 7, we consider situations in which preferences are observed only over acts defined in terms of finite state-spaces. We show how the DM's revealed probabilistic beliefs can still be identified from his preferences over multi-valued acts, under the assumption that his preferences over acts defined in terms of an appropriate larger state space are utility sophisticated; while not observed de facto, the larger preference relation is assumed to be observable in principle. All proofs are contained in the appendix.

2. BACKGROUND: IMPRECISE QUALITATIVE PROBABILITY

2.1. Precise Qualitative Probability: Savage's Theorem

A decision maker's probabilistic beliefs shall be modelled in terms of a partial ordering $\trianglerighteq$ on a σ-algebra of events Σ in a state space Θ, his "comparative likelihood relation", with the instance $A \trianglerighteq B$ denoting the DM's judgment that A is at least as likely as B. We shall denote the symmetric component of $\trianglerighteq$ ("is as likely as") by $\equiv$. The comparative likelihood relation can be viewed as representing a non-exhaustive (sub-)set of probabilistic judgments a DM is committed to, his *salient unambiguous beliefs*; these judgments, in turn, may reflect probabilistic information available to and accepted by him. The DM may have further "non-salient" probabilistic judgments not listed in $\trianglerighteq$; these will be inferred from his preferences in section 3.

We shall treat the comparative likelihood relation $\trianglerighteq$ as a non-behavioral primitive, which later will play the role of a constraint on and partial determinant of the DM's preferences. From a first-person point of view, this makes compelling intuitive sense. From a third-person point of view, the DM's salient unambiguous beliefs can be interpreted, for example, as accessible through direct communication rather than behavioral observation, or ascribed to him based on a shared understanding of the situation. While introducing such a non-behavioral primitive runs against certain strands of the "revealed preference" tradition in economics and decision theory, it needs to be pointed out that any Savage-style approach in which "behavior" is described in terms of preferences over mappings from states to consequences contains non-behavioral entities at its core: the state space, and the description of acts in terms of it. Philosophically, it is at least controversial whether a thoroughgoing revealed-preference approach is really possible; [11] moreover, in our everyday social reasoning and in the practice of economic modelling, direct ascriptions of beliefs are pervasive.

The following axioms are canonical; disjoint union is denoted by "+".

Axiom 1 (Partial Order)[12] *$\trianglerighteq$ is transitive and reflexive.*

Axiom 2 (Nondegeneracy) $\Theta \triangleright \emptyset$.

Axiom 3 (Nonnegativity) *$A \trianglerighteq \emptyset$ for all $A \in \Sigma$.*

Axiom 4 (Additivity) *$A \trianglerighteq B$ if and only if $A + C \trianglerighteq B + C$, for any C such that $A \cap C = B \cap C = \emptyset$.*[13]

Additivity is the hallmark of comparative *likelihood*. Normatively, it can be read as saying that in comparing two events in terms of likelihood, only states that are not common can matter.

For the sake of comparison, we briefly review first the case when the comparative likelihood relation is complete. It is well-known that, on finite state-spaces, Additivity is far from sufficient to guarantee the existence of a probability-measure representing the complete comparative likelihood relation; see Kraft-Pratt-Seidenberg (1959). On the other hand, a central part of Savage's famous characterization of SEU maximization was the demonstration that additivity suffices in the absence of probability "atoms". More specifically, Savage (1954) characterized the existence of convex-ranged probability measures; the probability measure π is **convex-ranged** if, for any event A and any $\alpha \in (0, 1)$, there exists an event $B \subseteq A$ such that $\pi(B) = \alpha\pi(A)$. We state a version of his result for the sake of further comparison. It requires one more axiom.

Axiom 5 (Partitioning) *For all A, B such that $A \triangleright B$ there exists a finite partition of Θ $\{C_1, ..., C_n\}$ such that $A \triangleright B \cup C_i$ for all $i \leq n$.*

Theorem 1 (Savage) *$\trianglerighteq$ is complete and satisfies Axioms 1 through 5 if and only if there exists a (unique) finitely additive, convex-ranged probability measure π on Σ such that for all $A, B \in \Sigma$:*

$$A \trianglerighteq B \text{ if and only if } \pi(A) \geq \pi(B).$$

[11] Cf. the contributions by Quine, Davidson, Nozick and Putnam for example. It is our impression that mainstream of contemporary analytical philosohpy would *deny* the possibility of a radical behavioral approach for the purposes of the social sciences.

[12] Technically, the proper label would be "preorder".

[13] In this notation, we quantify over all C disjoint from A and B.

An important feature of Savage's result is the uniqueness of the representing probability. It justifies the view that the comparative likelihood relation represents the DM's beliefs fully. This is non-trivial, and holds only rarely in finite state-spaces.

2.2 Dropping Completeness

It is a non-trivial task to generalize Savage's Theorem to the incomplete case. Here, the natural goal is a representation in terms of a set of finitely additive probability measures $\Pi \subseteq \Delta(\Theta)$ of the following form. For all $A, B \in \Sigma$:

$$A \trianglerighteq B \text{ if and only if } \pi(A) \geq \pi(B) \text{ for all } \pi \in \Pi. \tag{3}$$

Analogous representations for preferences over multi-valued acts have been given in particular by Bewley (1986) and Walley (1991) building on the work of Smith (1961). A comparative likelihood relation will be called **coherent** if it has a multi-prior representation (3). Note that if $\trianglerighteq$ satisfies (3) for some Π, then it satisfies (3) also for the convex hull of Π, as it does for the closure of Π (in the product or "weak*"-topology which will be assumed throughout). Thus, it is without loss of generality to assume Π to be a closed convex set, and uniqueness is understood to be at issue only within this class (which shall be denoted by $\mathcal{K}(\Delta(\Sigma))$.

Conceptually, coherence of a comparative likelihood relation can be interpreted as closure under inferences from the logic of probability. For example, additivity implies that if $A \trianglerighteq B$ then not $A^c \triangleright B^c$. If $\trianglerighteq$ is complete, one can infer that $B^c \trianglerighteq A^c$. This must hold under a multi-prior representation but does not follow automatically from additivity in the absence of completeness. To obtain coherence, we will rely on the following new axiom called "Splitting".

Axiom 6 (Splitting) *If* $A_1 + A_2 \trianglerighteq B_1 + B_2$, $A_1 \equiv A_2$ *and* $B_1 \equiv B_2$, *then* $A_i \trianglerighteq B_j$.

In words: If two events are split into two equally likely parts, then any part of the more likely event must be more likely than any part of the less likely event. Note that, as above, from additivity one can only infer that not $B_j \triangleright A_i$.

In finite settings, incomplete comparative likelihood relations will admit a *unique* multi-prior representation only in degenerate cases. To see this, note that any instance $A \trianglerighteq B$ constrains any $\pi \in \Pi$ to satisfy the condition $\pi(A) \geq \pi(B)$, which may also be written as $\pi \cdot (1_A - 1_B) \geq 0$. Thus, comparative likelihood relations can separate convex sets Π only through vectors taking values in $\{-1, 0, +1\}$ only. Thus, it is clear that, in general, many different convex sets will induce the same comparative likelihood relation. However, the uniqueness problem can be overcome if the comparative likelihood relation is "minimally complete" in an appropriate sense. Mathematically, what is needed is a sufficient supply of equality constraints on Π; it turns out that the following notion of "range convexity" is sufficient.

Definition 1 *A set of priors* Π *is* **convex-ranged** *if, for any event* A *and any* $\alpha \in (0,1)$, *there exists an event* $B \subseteq A$ *such that* $\pi(B) = \alpha\pi(A)$ *for all* $\pi \in \Pi$.

Note that while range convexity of Π implies range convexity of every $\pi \in \Pi$, the converse is far from true. Range convexity has the following axiomatic counterpart which assumes that any event can be split into two equally likely parts.

Axiom 7 (Equidivisibility) *For any* $A \in \Sigma$, *there exists* $B \subseteq A$ *such that* $B \equiv A \backslash B$.

Note also that Equidivisibility implies that any atom A of Σ must be a null-event.[14] In the following, we will often refer to an comparative likelihood relation satisfying Equidivisibility as **minimally complete**.

Finally, Savage's Partitioning axiom in its role as a continuity condition is no longer adequate. With incompleteness, it is no longer sufficient nor even necessary. Indeed, even on independent conceptual grounds, it seems desirable to express the notion of "continuity in probability" in terms of a pure condition such as the following one which is applicable to any qualitative probability, whether precise or not, and to any state space. It relies on the following notion of a $\frac{1}{K}-$event: A is a $\frac{1}{K}-$**event** if there exist at $K-1$ mutually disjoint events A_i, disjoint from A, such that $A \trianglelefteq A_i$ for all i. Clearly, for coherent $\trianglerighteq$ and any $\pi \in \Pi$ and any $\frac{1}{K}-$event A, $\pi(A) \leq \frac{1}{K}$; if Π is convex-ranged, the converse holds as well.

Axiom 8 (Continuity) *If not* $A \trianglerighteq B$, *then there exists* $K < \infty$ *such that, for any* $\frac{1}{K}-$ *events* C, D, *it is not the case that* $A \cup C \trianglerighteq B \backslash D$.

Note that continuity is entailed by coherence.

Theorem 2 *A relation* $\trianglerighteq$ *has a multi-prior representation with a convex-ranged set of priors* Π *if and only if it satisfies Partial Order, Additivity, Nonnegativity, Splitting, Continuity, Equidivisibility, and Nondgeneracy.*

The representing Π *is unique in* $\mathcal{K}(\Delta(\Sigma))$.

[14] That is, for any A for which $B \subset A$ implies $B = \varnothing$, Equivdisibility implies $A \equiv \emptyset$.

In view of Theorem 2, we shall refer to a comparative likelihood relation with a convex-ranged multi-prior representation as an **imprecise qualitative probability** (relation). We shall sketch the idea of the proof of Theorem 2 with a bit of "reverse engineering". While its construction is crucial to understanding the logic behind most of the results to follow, the results themselves can be understood without it. Hence the following passage may be skimmed on a first reading.

One can extend every imprecise qualitative probability represented by the convex-ranged set of priors Π to the set $B(\Sigma, [0,1])$ of finite-valued functions $Z : \Theta \to [0,1]$ by associating with each Z an equivalence class $[Z]$ of events $A \in \Sigma$ as follows. Let $A \in [Z]$ if, for some appropriate partition of Θ $\{E_i\}$, $Z = \sum z_i 1_{E_i}$, and, for all $i \in I$ and $\pi \in \Pi : \pi(A \cap E_i) = z_i \pi(E_i)$. It is easily seen that for any two $A, B \in [Z] : \pi(A) = \pi(B)$ for all $\pi \in \Pi$, and thus $A \equiv B$. One therefore arrives at a well-defined partial ordering on $B(\Sigma, [0,1])$, denoted by $\widehat{\trianglerighteq}$, by setting $Y \widehat{\trianglerighteq} Z$ if $A \trianglerighteq B$ for some $A \in [Y]$ and $B \in [Z]$. It is easily verified that this ordering is monotone, continuous and satisfies the independence axiom, i.e. that

$$Y \widehat{\trianglerighteq} Z \text{ if and only if } \alpha Y + (1-\alpha)X \widehat{\trianglerighteq} \alpha Z + (1-\alpha)X \text{ for any } X, Y, Z \text{ and } \alpha \in (0,1].$$

The proof of Theorem 2 proceeds by constructing this extension from the qualitative probability relation and by deriving the three properties of the induced relation from the axioms on the primitive relation. Deriving the independence axiom entails the hardest work. It can be viewed as consisting of the following additive and a multiplicative invariance conditions:

$$Y \widehat{\trianglerighteq} Z \text{ if and only if } Y + X \widehat{\trianglerighteq} Z + X \text{ for any } X, Y, Z \ ,$$

and

$$Y \widehat{\trianglerighteq} Z \text{ if and only if } \alpha Y \widehat{\trianglerighteq} \alpha Z \text{ for any } Y, Z \text{ and } \alpha \in (0,1].$$

The two conditions correspond to the additive and splitting axioms characteristic of an imprecise qualitative probability, respectively. The proof then invokes a version of Bewley's (1986) Theorem to obtain the desired multi-prior representation.

In view of the above construction, an imprecise qualitative probability can be viewed as inhabiting a mixture-space in the manner of Anscombe-Aumann without reference to an "extraneous" randomization device. On the other hand, state-spaces with a continuous randomization device furnish an important example of minimally complete imprecise qualitative probabilities which we shall refer to as the Anscombe-Aumann example. Specifically, consider a state space that can be written as $\Theta = \Theta_1 \times \Theta_2$, where the (finite) space Θ_1 is the space of "generic states" , and Θ_2 that of independent "random states" with associated $\sigma-$algebra Σ_2. The "continuity" and stochastic independence of the random device are captured by an imprecise qualitative probability $\trianglerighteq_{AA}$ defined on $\Sigma = 2^{\Theta_1} \times \Sigma_2$ that satisfies the following two conditions.

AA1) The restriction of $\trianglerighteq_{AA}$ to $\{\Theta_1\} \times \Sigma_2$ is complete and satisfies Partitioning.

AA2) $\sum_{\theta_1} \{\theta_1\} \times T_{\theta_1} \trianglerighteq_{AA} \sum_{\theta_1} \{\theta_1\} \times T'_{\theta_1}$ if and only if, for all $\theta_1 \in \Theta_1$,

$\Theta_1 \times T_{\theta_1} \trianglerighteq_{AA} \Theta_1 \times T'_{\theta_1}$.

While the first condition ensures the existence of a convex-ranged probability measure over random events $\pi_2 \in \Delta(\Sigma_2)$, the second describes their stochastic independence. Consider any imprecise qualitative probability $\trianglerighteq$ containing $\trianglerighteq_{AA}$. By AA1 and AA2, $\trianglerighteq_{AA}$ and thus $\trianglerighteq$ satisfy Equidivisibility. From Theorem 2, one obtains the existence of a unique closed convex set of finitely additive probability measures Π representing $\trianglerighteq$; it has the property that, for all $\pi \in \Pi$ and all $\theta_1 \in \Theta_1$, $\text{marg}_{\Sigma_2}\pi(./\theta_1) = \pi_2$, reflecting the stochastic independence of the random device.[15]

3. BETTING PREFERENCES AND PROBABILISTIC BELIEFS

3.1. Compatibility of Preferences with Probabilistic Beliefs

Consider now a DM described by a preference ordering over acts and salient beliefs over events. Let X be a set of *consequences*. An *act* is a finite-valued mapping from states to consequences, $f : \Theta \to X$, that is measurable with respect to a σ-algebra of events Σ; the set of all acts is denoted by $\mathcal{F}$. A *preference ordering* $\succsim$ is a weak order (complete and transitive relation) on $\mathcal{F}$. We shall write $[x_1, A_1; x_2, A_2; ...]$ for the act with consequence x_i in event A_i; constant acts $[x, \Theta]$ are typically referred to by their constant consequence x.

The DM also has salient beliefs described by a coherent comparative likelihood relation $\trianglerighteq^0$ on Σ. This relation represents *some* of the DM's probabilistic beliefs; he may have others not included in it. Later in this section we shall define a richer relation $\trianglerighteq^*$ derived from his preferences that captures *all* probabilistic beliefs that can be meaningfully described to the DM. Thus, the datum for the following is a pair $(\succsim, \trianglerighteq^0)$. We shall ask in section 5 under what conditions $\trianglerighteq^0$ becomes

[15] For an earlier representation of the Anscombe-Aumann framework in a Savage setting, see Klibanoff (2001a). A different route to "subjectivizing" the Anscombe-Aumann framework based on a rich set of consequences rather than states is presented in Ghirardato et al. (2001d).

redundant and can be replaced by $\trianglerighteq^*$ so that $\trianglerighteq^*$ itself is purely definable in preferences, and a "fully behavioral" account is possible.

Salient beliefs determine primarily preferences over bets. A *bet* is a pair of acts with the same two outcomes, i.e. a pair of the form $([x, A; y, A^c], [x, B; y, B^c])$. Fundamental is the following rationality requirement on the relation between preferences and probabilistic beliefs.

Axiom 9 (Compatibility) *For all $A, B \in \Sigma$ and $x, y \in X$ such that $x \succ y$:*

$$[x, A; y, A^c] \succsim [x, B; y, B^c] \text{ if } A \trianglerighteq^0 B, \text{ and}$$
$$[x, A; y, A^c] \succ [x, B; y, B^c] \text{ if } A \triangleright^0 B.$$

Note that, since the relation $\trianglerighteq^0$ is required to be coherent, Compatibility builds in respect for all inferences from the "logic of probability". Applied to Example 1 of the Introduction, for example, axiom 9 delivers restriction (2) on preferences involving ambiguous events.[16] Compatibility becomes especially powerful when salient beliefs are minimally complete, hence an imprecise qualitative probability. In this case, the associated preference relation will be referred to as "**minimally unambiguous**".

3.2 Regular Preferences over Bets

As is customary, we shall assume that preferences over bets depend only on the events involved, not on the stakes. This is captured by Savage's axiom P4. In the context of ambiguity, P4 is still a natural "regularity" assumption , but it seems doubtful that it can still be viewed as a rationality axiom.[17]

Axiom 10 (Consistent Confidence Ranking, P4)

For all $x, y, x', y' \in \mathcal{X}$ such that $x \succ y$ and $x' \succ y'$ and all $A, B \in \Sigma$:

$[x, A; y, A^c] \succsim [x, B; y, B^c]$ *iff* $[x', A; y', A^c] \succsim [x', B; y', B^c]$.

[16]Note also that by compatibility, the betting preferences must reflect the beliefs attributed to the decision maker. It is not clear what further, *purely behavioral* evidence could be brought to bear on the question whether the DM "in fact" has those beliefs. On the other hand, as discussed in section 5.2, compatiblity with a hypothesized imprecise qualitative probability by itself does not seem warrant imputing that relation as probabilistic beliefs to the decision maker.

[17]Note that P4 can be *deduced* from the assumption that the preference relation be compatible with some *complete* qualitative probability. This justification of P4, however, evidently eliminates any room for ambiguity.

Given this axiom, one can define an induced **confidence relation**[18] $\geq$ (with symmetric component $\doteqdot$) on Σ by setting

$$A \geq B \text{ iff } [x, A; y, A^c] \succsim [x, B; y, B^c] \text{ for some } x \succ y.$$

In the following, we will typically refer to betting preferences through their associated confidence relation; note that compatibility with salient beliefs can be rewritten as the requirement that $\geq$ contain $\trianglerighteq^0$ as a relation.

To ensure real-valued representability, betting preferences are assumed to satisfy the following "Archimedean" property. We will say that A' is a $\frac{1}{n}$**−fraction** of A if there is a partition of A $\{A', A_2, .., A_n\}$ such that $A' \equiv^0 A_i$ for all $i \geq 2$.

Axiom 11 (Archimedean) *If $A > B$, there exists $K < \infty$ and $\frac{1}{K}$−fractions C of A^c and D of B^c such that $A \backslash C > B$ and $A > B + D$.*

Two well-behavedness assumptions are also needed; below in Proposition 3 of section 4.2.3., they will be derived from P4 via a rationality argument. Here we shall motivate them on their own somewhat loosely.

Axiom 12 (Union Invariance) *For any T that is a $\frac{1}{n}$−fraction of Θ, and for any $A, B \in \Sigma$ such that $A \cap T = B \cap T = 0$: $A \geq B$ if and only if $A + T \geq B + T$.*

The essence of the requirement on T is its unambiguity in terms of the basic $\trianglerighteq^0$; conceptually, it is immaterial that T be assigned a probability of the form $\frac{1}{n}$. The axiom is intuitive, as it says that comparative confidence is not affected by the addition of events of known probability. Indeed, it is sufficiently intuitive for Epstein-Zhang (2001) to turn it around and, by formulating it as a condition on an event T, to make it their linchpin for the very definition of an event T as "revealed unambiguous".

While Union Invariance can be viewed as a restricted version of Additivity to instances in which the addendum is unambiguous, the following axiom analogously requires invariance to splits when these are unambiguous.

Axiom 13 (Splitting Invariance) *For any A, B and $\frac{1}{n}$-fractions A', B' of A and B respectively, $A \geq B$ if and only if $A' \geq B'$.*

[18] In contrast to our terminology, $\geq$ is often referred to as the DM's "revealed likelihood" relation, following Savage (1954). In the context of ambiguity, this terminology seems however inappropriate, as the relation $\geq$ reflects not only a DM's beliefs but also his "ambiguity attitude".

Preference relations (as well as the associated confidence relation $\trianglerighteq$) shall be called **regular** if they satisfy the P4, Continuity, Union Invariance, Splitting Invariance, and the maintained assumption of compatibility with the salient imprecise qualitative probability $\trianglerighteq^0$.

For any $\trianglerighteq^0$-non-null $A \in \Sigma$, let Λ_A^0 denote the family of **salient unambiguous events conditional on** A given as the set of all $B \subseteq A$ such that

$$B \in \Lambda_A^0 \text{ iff there exists } \alpha \in (0,1] \text{ such that } \pi(B) = \alpha\pi(A) \text{ for all } \pi \in \Pi_0,$$

and let $\pi^0(B/A)$ denote the unique α satisfying this condition; $\pi^0(B/A)$ represents the unambiguous conditional probability of B given A. In the case of $A = \Theta$, we shall write simply Λ^0 and $\pi^0(B)$.

It is easily verified that an Archimedean, minimally unambiguous confidence relation has a unique representation in terms of a non-additive set function $\rho : \Sigma \to [0,1]$ such that

$$A \geq B \text{ if and only if } \rho(A) \geq \rho(B),$$

with $\rho(\Theta) = 1$ and such that ρ is additive on the family of salient unambiguous events Λ^0; ρ will be call the DM's **confidence measure.** Confidence is calibrated in terms of equivalent unambiguous probability: $\rho(A) = \pi^0(B)$ for any $B \in \Lambda^0$ such that $B \doteqdot A$. A confidence relation $\geq$ that is compatible with the imprecise qualitative probability $\trianglerighteq^0$ induces a well-defined $\widehat{\geq}$ on the mixture-space $B(\Sigma,[0,1])$ associated with $\trianglerighteq^0$; $\widehat{\geq}$ is defined as follows.

$$Y \widehat{\geq} Z \text{ if } A \geq B, \text{ for any } A \in [Y] \text{ and } B \in [Z],$$

using the [ ]-notation of section 2.2. To verify well-behavedness, simply observe that by construction of the mixture space, for any two $A, B \in [Z]$, one has $A \equiv^0 B$, hence by compatibility also $A \doteqdot B$. Let $\widehat{\rho}$ denote the associated unique extension of ρ to $B(\Sigma,[0,1])$ given by

$$\widehat{\rho}(Y) = \rho(A) \text{ for any } A \in [Y].$$

Again, by the construction of the mixture-space, this is well-defined, and by construction

$$Y \widehat{\geq} Z \text{ if and only if } \widehat{\rho}(Y) \geq \widehat{\rho}(Z).$$

Note that $\widehat{\rho}(c\mathbf{1}) = c$ for any $c \in [0,1]$; moreover, $\widehat{\rho}$ is *monotone* ($\widehat{\rho}(Y) \geq \widehat{\rho}(Z)$ whenever $Y \geq Z$) by compatibility of $\geq$ with $\trianglerighteq^0$. $\widehat{\rho}$ is *c-additive* if $\widehat{\rho}(Y + c\mathbf{1}) = \widehat{\rho}(Y) + c$; $\widehat{\rho}$ is *positively homogeneous* if $\widehat{\rho}(\alpha Y) = \alpha\widehat{\rho}(Y)$ for any $\alpha \in [0,1]$; $\widehat{\rho}$ is *c-linear* if it is c-additive and positively homogeneous. If $\widehat{\rho}$ is *c-linear*, in view of its monotonicity it is also continuous in the sup-norm topology.

We note the following properties of $\widehat{\rho}$.

Fact 1 *$\widehat{\rho}$ is c-additive (resp. positively homogeneous) if and only if $\geq$ satisfies Union Invariance (resp. Splitting Invariance).*

Thus, for preferences over bets, regularity is the analogue in the present framework to the "certainty independence" axiom (in the Anscombe-Aumann framework) due to Gilboa-Schmeidler (1989).

3.3 Revealed Unambiguous Beliefs

One can now turn the compatibility question around and ask, for given preferences over bets $\succsim$, which imprecise qualitative probability relation $\trianglerighteq$ the preferences are compatible with. This question is of interest from a third-person point of view, as it asks which probabilistic beliefs it is meaningful to *impute* to the observed DM. The following results show that if betting preferences are regular, there is always a unique largest imprecise qualitative probability with which the preference relation is compatible and that extends the salient belief relation $\trianglerighteq^0$. It thus makes conceptual sense to identify this maximal relation as describing *the* unambiguous beliefs *revealed* by the DM.

Theorem 3 *1. Let $\geq$ be a regular confidence relation compatible with the imprecise qualitative probability $\trianglerighteq^0$. Then there exists a maximal imprecise qualitative probability $\trianglerighteq^*$ such that $\geq \supseteq \trianglerighteq^* \supseteq \trianglerighteq^0$.*

2. $\trianglerighteq^*$ is given by

$$A \trianglerighteq^* B \text{ if and only if } A' + C \geq B' + C,$$

for all C and all A', B' such that, for some $n < \infty$, A' is a $\frac{1}{n}-$fraction of A and B' is a $\frac{1}{n}-$fraction of B.

3. $\Pi_{\trianglerighteq^*}$ is the smallest closed, convex set $\Pi \subseteq \Pi_{\trianglerighteq^0}$ such that for all A, B such that $A > B$, there exists $\pi \in \Pi$ such that $\pi(A) > \pi(B)$.

In the future, we shall write Π^* for $\Pi_{\trianglerighteq^*}$, and denote by Λ_A^* the family of revealed unambiguous events conditional on A defined analogously to Λ_A^0, with associated unambiguous conditional probability measure $\pi^*(./A)$.

Remark 1. The characterizing condition in part 2) is "designed" to build in the Additivity and Splitting properties of an imprecise qualitative probability. The latter is essential; if one would stipulate as the criterion for unambiguous preference of A over B merely that $A + C \geq B + C$ for all C, then A would be unambiguously equally likely to A^c whenever there is equal confidence in A than in A^c (i.e. $A \doteqdot A^c$), simply because there is no non-empty set disjoint from both A and A^c; evidently, this makes little sense.

Remark 2. The "price" of needing to build in the Splitting Axiom is the need to rely on the salient belief relation $\trianglerighteq^0$ relative to which unambiguous splits are defined. It may be the case that in part 1) of the Theorem, one could weaken the requirement that $\trianglerighteq^*$ contain the relation $\trianglerighteq^0$ to that of being minimally complete; indeed, we expect this to be frequently be the case. This would have the advantage of characterizing $\trianglerighteq^*$ in terms of preferences only[19]. On the other hand, it is easy to give examples in which $\trianglerighteq$ is compatible with coherent comparative likelihood relations $\trianglerighteq'$ that are not minimally complete and that are not contained in $\trianglerighteq^*$; in those cases, there is then no unique maximal coherent comparative likelihood relation. In section 5, we shall show that one can overcome the need for a reference relation $\trianglerighteq^0$ if one is willing to make substantial rationality assumptions on the DM's preferences over multi-valued acts.

Remark 3. The idea of using a maximal "independent" subrelation to define revealed unambiguous beliefs was first proposed in Nehring (1996) which presented an analogue to Theorem 3 in an Anscombe-Aumann framework; see section 7 for further discussion.

Remark 4. Regularity of preferences has been defined relative to the salient belief relation $\trianglerighteq^0$. Since that relation is understood to describe the DM's beliefs non-exhaustively, there is clearly a degree of arbitrariness in specifying a DM's salient beliefs, as further beliefs might also have been deemed salient. Thus, to be conceptually well-defined, the regularity axioms Union and Splitting Invariance should apply equally to any imprecise qualitative probability $\trianglerighteq$ such that $\trianglerighteq^0 \subseteq \trianglerighteq \subseteq \trianglerighteq^*$. For Splitting Invariance, this follows[20] from the following Fact which is a straightforward consequence of the range convexity of Π^0 and the inclusion of Π^* in Π^0.

Fact 2 *For any $A, B \in \Sigma$ such that $A \in \Lambda_B^*$, there exist $A' \in \Lambda_B^0$ such that $\pi^*(A/B) = \pi^*(A'/B) = \pi^0(A'/B)$.*

[19] With the big qualification that the imputation of a given state space is a major non-behavioral element in any Savage-style analysis.

[20] To verify this, consider $A, B \in \Sigma$ and $\frac{1}{n}$−fractions (with respect to $\trianglerighteq^*$) A', B'. By Fact 2, there exist $\frac{1}{n}$−fractions with respect to $\trianglerighteq^0$ A'', B'' such that $A'' \equiv^0 A'$ and $B'' \equiv^0 B'$. By Compatibility, $A' \doteqdot A''$ and $B' \doteqdot B''$, whence $A' \geq B'$ iff $A'' \geq B''$. Thus Splitting Invariance with respect to $\trianglerighteq^*$ is equivalent to Splitting Invariance with respect to $\trianglerighteq^0$.

For Union Invariance, this can likewise be derived from Fact 2 using Splitting Invariance, but it also follows directly from Proposition 1 below. Finally, the Archimedean axiom is evidently weaker relative to $\trianglerighteq^*$ than to $\trianglerighteq^0$.

3.4 Revealed Unambiguous Events

In the preceding literature, attention has been paid not to unambiguity of preferences or beliefs in general, but to unambiguity of events (see in particular Ghirardato-Marinacci (2001a), Nehring (1999), Epstein-Zhang (2001)). Presumably, an event is "revealed unambiguous" if the DM reveals an unambiguous probabilistic belief concerning its occurrence. This is captured by the following formal definition.

Definition 2 *The event A is* **revealed unambiguous** *if $A \equiv^* T$ for some $T \in \Lambda^0$.*

Revealed unambiguous events can be characterized in the following equivalent ways.

Proposition 1 *For a regular confidence relation $\geq$, the following statements are equivalent:*

1. *A is revealed unambiguous.*
2. *For all $\pi, \pi' \in \Pi^* : \pi(A) = \pi'(A)$.*
3. *For all $A' \in \Lambda^0_A$ and $B \in \Sigma : \rho(A' + B) - \rho(B) = \rho(A')$.*
4. *For all $A' \in \Lambda^0_A$ and $B, C \in \Sigma : B \geq C$ iff $B + A' \geq C + A'$.*

Let Λ^* denote the family of all revealed unambiguous events (derived from $\geq$). By part ii), Λ^* is a $\lambda-$**system**, i.e. it contains Θ and is closed under complementation and (finite) disjoint union. Moreover, the confidence measure ρ is additive on Λ^* and can be extended to any $\pi \in \Pi^*$. In particular, ρ is a qualitative probability iff $\Lambda^* = \Sigma$.

Proposition 1 allows a direct comparison to Epstein-Zhang's (2001) definition, as it applies to preferences over bets. It is formulated here for the sake of comparability in terms of the confidence relation $\trianglerighteq$: an event A is **EZ-unambiguous** if, for any B, C disjoint from $A : B \geq C$ iff $B + A \geq C + A$, and if the same holds when A is replaced by A^c.[21] In view of part (4) of Proposition 1, an

[21]Epstein-Zhang (2001) definition for multi-valued acts builds in probabilistic sophistication over unambiguous acts; we obtain this as an implication of a separate rationality axiom on preferences over multi-valued acts called "Revealed Stochastic Dominance"; see section 4.1.

event is revealed unambiguous in our sense if and only if all of its fractions are EZ-unambiguous; the Epstein-Zhang definition is thus the weaker, more permissive one.[22] Epstein-Zhang (2001) show (under by and large weaker assumptions than regularity), that the family of EZ-unambiguous events is a λ-system, and that ρ is additive on this family. Thus, it might appear that EZ-unambiguous events can be understood as events over which one can attribute probabilistic beliefs to the DM. But such an interpretation would not valid in general. In particular, as noted in both Epstein (1999) and Nehring (1999), additivity on a λ-system does not imply extendability to a probability measure on all of Σ; in other words, the comparative likelihood relation associated with the additive measure on the family of EZ-unambiguous events may be *inconsistent,* and thus not interpretable as corresponding to any well-defined set of probabilistic beliefs.

One can extend Definition 2 to say that the event A is **revealed unambiguous conditional on** B (with B non-null) if $A \equiv^* B'$ for some $B' \in \Lambda_B^0$. Analogously to Proposition 1, the event A is revealed unambiguous conditional on B if and only if, for all $\pi, \pi' \in \Pi^*$ such that $\pi(B) > 0$ and $\pi'(B) > 0$, $\pi' : \pi(A/B) = \pi'(A/B)$. However, the third and fourth characterizations do not seem to generalize straightforwardly.

3.5. Uniform Caution

Theorem 3 suggests a natural decomposition of the "confidence" in an event A in a belief component given by the event A's lower and upper probabilities revealed by the confidence relation, and a psychological "ambiguity reaction" describing the extent to which "confidence" is determined by the most pessimistic respectively the most optimistic probability assignment compatible with the confidence relation. After defining formally the DM's ambiguity reaction in terms of an event-specific "degree of caution", we shall characterize the model that results from assuming constancy of the degree of caution across events. This section is an adaptation of section 4 of Ghirardato-Maccheroni-Marinacci (2001c).

For any event $A \in \Sigma$, define the event's revealed lower probability $\pi^*_{\min}(A) := \min_{\pi \in \Pi^*} \pi(A)$ and upper probability $\pi^*_{\max}(A) = \max_{\pi \in \Pi^*} \pi(A)$. Evidently, an event is revealed ambiguous iff $\pi^*_{\min}(A) < \pi^*_{\max}(A)$. Compatibility implies that $\pi^*_{\min}(A) \leq \rho(A) \leq \pi^*_{\max}(A)$. Hence, for any revealed

[22] In the special case of ambiguity averse betting preferences in the sense of section 6, however, "unambiguous" and EZ-unambiguous events coincide.

ambiguous event $A \in \Sigma \backslash \Pi^*$, there exists a unique number $\gamma (=: \gamma(A)) \in [0,1]$ such that

$$\rho(A) = \gamma(A)\pi^*_{\min}(A) + (1-\gamma(A))\pi^*_{\max}(A).$$

$\gamma(A)$ is the "degree of caution" exercised in assessing a bet on A. The maximal degree of caution $\gamma(A) = 1$ describes exclusive reliance on the lower probability in this assessment; likewise, the minimal degree of caution (that is, the maximal degree of "boldness") $\gamma(A) = 0$ describes exclusive reliance on the upper probability.[23]

Empirically, there is substantial evidence that $\gamma(A)$ depends on the event A; see for example Wakker (2001). Nonetheless, for applications it is clearly of interest to focus on the special case in which this event-dependence does not happen, i.e. in which $\gamma(A) = \gamma$ for all A. In this case, γ is plausibly interpreted as representing the DM's event-independent overall "degree of caution". This case will be referred to as the "**uniform model**".[24] Note that if the degree of caution is constant across events, confidence in an event is monotone in the probability interval of the event, and in this sense determined by the unambiguous beliefs about the event:

$$\rho(A) \geq \rho(B) \text{ whenever } \pi^*_{\min}(A) \geq \pi^*_{\min}(B) \text{ and } \pi^*_{\max}(A) \geq \pi^*_{\max}(B).$$

Translated into the following condition on preferences, this yields a characterization of the uniform model.

Axiom 14 (Interval Monotonicity) *Suppose that, given $A, B \in \Sigma$ it is the case that, for every $T \in \Lambda^*$, $A \trianglelefteq^* T$ implies $B \trianglelefteq^* T$, and that $B \trianglerighteq^* T$ implies $A \trianglerighteq^* T$. Then $A \geq B$.*

Proposition 2 *Let $\geq$ be a regular confidence relation compatible with the salient imprecise qualitative probability $\trianglerighteq^0$. Then $\geq$ has the representation*

$$\rho(A) = \gamma\pi^*_{\min}(A) + (1-\gamma)\pi^*_{\max}(A), \text{ for all } A \in \Sigma,$$

if and only if it satisfies Interval Monotonicity.

[23]It would be inappropriate to refer to $\gamma(A)$ as a "degree of pessimism", since this would suggest a statement about the DM's beliefs, but these are already described by the probability interval $[\pi^*_{\min}(A), \pi^*_{\max}(A)]$.

[24]Note that the regularity axioms Union and Splitting Invariance can be interpreted as invariance conditions on $\gamma(A)$. Indeed, the former is equivalent to the condition

$$\text{For any } A \in \Sigma\backslash\Lambda^0 \text{ and any } T \in \Lambda^0 \text{ such that } A \cap T = \emptyset : \gamma(A+T) = \gamma(A),$$

while the latter corresponds to the condition

$$\text{For any } A \in \Sigma\backslash\Lambda^0 \text{ and any } B \in \Lambda^0_A : \gamma(B) = \gamma(A).$$

Due to the mixture-space representation of minimally unambiguous confidence relations, Proposition 2 is a direct consequence of Theorem 12 of Ghirardato et al. (2001c) which characterizes the interval expected utility model defined in Fact 4 below (called by them the α-minimum expected utility model) in an Anscombe-Aumann framework.[25]

4. PREFERENCES OVER MULTI-VALUED ACTS

4.1 Revealed Stochastic Dominance

In this section, we shall introduce two rationality principles that determine how preferences over bets constrain preferences over general acts. A natural minimal rationality requirement under ambiguity is respect for stochastic dominance, to the extent that this criterion can be applied on the basis of the DM's revealed unambiguous beliefs. Thus, say that f **revealed stochastically dominates** g if, for all $x \in X$, $\{\theta \mid f(\theta) \succsim x\} \trianglerighteq^* \{\theta \mid g(\theta) \succsim x\}$.

Axiom 15 (Revealed Stochastic Dominance) *$f \succsim g$ whenever f revealed stochastically dominates g.*

An act f is *unambiguous* if, for all $x \in X$, $\{\theta \mid f(\theta) \succsim x\} \in \Lambda^*$. Letting π^* denote the restriction of any $\pi \in \Pi^*$ to Λ^*, satisfaction of Revealed Stochastic Dominance entails respect for ordinary stochastic dominance of outcomes with respect to π^* as described by the following condition. For any unambiguous acts f, g:

$$f \succsim g \text{ if } \pi^* (\{\theta \mid f(\theta) \succsim x\}) \geq \pi^* (\{\theta \mid g(\theta) \succsim x\}) \text{ for all } x \in X. \tag{4}$$

In particular, $f \sim g$ whenever f and g induce the same probability distribution over outcomes; the agent's preferences over unambiguous acts are therefore *probabilistically sophisticated* in the sense of Machina-Schmeidler (1992), but otherwise unconstrained, leaving plenty of room for Allais-type phenomena.[26]

[25] For $\gamma = 1$, this had already been shown by Nehring (1996).

[26] Obtaining probabilistic sophistication on unambiguous acts was a central desideratum of Epstein-Zhang (2001). While they achieved this goal as a consequence of their definition of an unambiguous event, it seems cleaner to separate such a definition from Revealed Stochastic Dominance as a rationality requirement, as done here. Besides this methodological difference, a further substantive difference to Epstein-Zhang comes from the fact that Revealed Stochastic Dominance entails significant restrictions also on comparisons of ambiguous acts, since it applies whenever the events of the form $\{\omega \mid f(\omega) \succsim x\}$ are unambiguously comparable across the two acts in terms of revealed

Generalizing equation (4), the following characterization of revealed stochastic dominance is straightforward from Theorem 3 and the well-known utility characterization of ordinary stochastic dominance.

Fact 3 *f revealed stochastically dominates g if and only if $E_\pi u \circ f \geq E_\pi u \circ g$ for all $\pi \in \Pi^*$ and all $u : X \rightarrow \mathbf{R}$ such that $u(x) \geq u(y)$ whenever $x \succsim y$.*

Fact 3 suggests that Revealed Stochastic Dominance is the strongest rationality requirement that relies on ordinal information about the valuation of consequences only. Revealed Stochastic Dominance has a precursor of sorts in the literature, Sarin-Wakker's (1992) "Cumulative Dominance" axiom. Say that f *cumulatively dominates* g if, for all $x \in X$, $\{\theta \mid f(\theta) \succsim x\} \geq \{\theta \mid g(\theta) \succsim x\}$; the preference relation satisfies Cumulative Dominance if $f \succeq g$ whenever f cumulatively dominates g. Thus, Cumulative Dominance is simply Revealed Stochastic Dominance modified to rely on arbitrary confidence comparisons rather than just on unambiguous ones. Conceptually, this changes the nature of the condition fundamentally, as the confidence relation $\geq$ cannot be interpreted as describing the DM's beliefs only. By consequence, Cumulative Dominance cannot be viewed as a rationality principle.[27] Instead, in Sarin-Wakker (1992), in effect it gives rise to a "procedural description" of how to construct CEU preferences over general multi-valued acts from the confidence relation and SEU preferences over unambiguous acts. Note that since CEU preferences satisfy Cumulative Dominance, they satisfy Revealed Stochastic Dominance a fortiori.

4.2 Utility Sophistication

4.2.1. Tradeoff Consistency.—

There are two basic types of departures from expected utility maximization: ambiguity, as illustrated by the Ellsberg paradox, and "probabilistic risk-aversion" (non-linear weighting of utilities), as shown by the Allais-paradox. It thus seems worthwhile to demarcate those cases in which all departures from expected utility maximization can be attributed to one of these two sources exclusively. "Probabilistic sophistication" (over all acts) is one such demarcation: it excludes Ellsbergian phenomena while preserving great flexibility with respect to phenomena of probabilistic risk-aversion. Since ambiguity concerns the very nature of uncertainty itself, it is arguably more basic a departure

likelihood, even if these events themselves are ambiguous. Speaking loosely, Revealed Stochastic Dominance implies for example that preferences over "almost unambiguous" acts are close to preferences over unambiguous acts proper.

[27]This has been pointed out before in Nehring (1994).

from the expected utility paradigm than probabilistic risk-aversion, and arguably harder to rule out on normative grounds. It seems therefore especially important to possess a model that allows to zoom in on ambiguity alone. In this subsection, we thus want to propose a complementary concept of "utility sophistication" which accommodates ambiguity in a very general way but excludes Allais-type phenomena systematically. While there are examples of utility sophisticated models in the literature (in particular the seminal contributions by Schmeidler (1989) and Gilboa-Schmeidler (1989) in an Anscombe-Aumann-framework), and while a "utility sophisticated" outlook has been advocated forcefully in Ghirardato-Marinacci (2001a), a general definition of "utility sophistication" has not yet been attempted, as far as we know.

To motivate the key axiom underlying utility sophistication, consider first the risk-neutral case in which consequences are given in amounts of income, and in which the DM ranks unambiguous acts according to expected income. Specifically, consider the DM's choice between the two acts $[x, A; y, B; f_{-A-B}]$ and $[x', A; y', B; f_{-A-B}]$, with A judged equally likely to B, that is: $A \equiv^0 B$. Note that, conditional on the event $A \cup B$, the DM needs to compare two fifty-fifty lotteries. Thus, there seems to be little doubt what our risk-neutral DM should do: choose the act with the higher conditional expected payoff; that is, he should prefer $[x, A; y, B; f_{-A-B}]$ over $[x', A; y', B; f_{-A-B}]$ weakly whenever $\frac{1}{2}x + \frac{1}{2}y \geq \frac{1}{2}x' + \frac{1}{2}y'$.

Next, abandon the assumption of risk-neutrality and consider choices among unambiguous acts with two outcomes, each of which has subjective probability one half. Specifically, consider a choice between $f = [x, A; y, A^c]$ and $g = [x', A; y', A^c]$ such that $x \succ x'$, $y' \succ y$ and $A \equiv^0 A^c$. According to a classical interpretation of expected utility theory, a DM ("You") should choose f over g exactly if You assess the utility gain from x over x' to exceed the loss of obtaining y rather than y'. Conversely, a preference of f over g is naturally interpreted as revealing a greater utility gain from x over x' than from y' over y. This intuition generalizes to choices of the form $[x, A; y, B; f_{-A-B}]$ versus $[x', A; y', B; f_{-A-B}]$ whenever the events A and B are judged equally likely $(A \equiv^0 B)$: for also in this more general case, the comparison of these utility gains is the only remaining, hence decisive factor in the choice. To compare the two acts, the DM simply does not need to consider his (possibly imprecise) assessment of the likelihood of the union $A + B$, nor the payoffs in states outside $A + B$. This motivates the following rationality axiom which requires that the DM's preferences must be rationalizable in terms of a comparison of utility differences that is consistent across choices of the above kind.

Axiom 16 (Tradeoff Consistency) *For all* $x, y, x', y' \in X$, $f, g \in \mathcal{F}$ *and* A, B, A', B' *such that* $A \cap B = A' \cap B' = \emptyset$ *and* $A \equiv^0 B$ *as well as* $A' \equiv^0 B'$: $[x, A; y, B; f_{-A-B}] \succsim [x', A; y', B; f_{-A-B}]$ *if and only if* $[x, A'; y, B'; g_{-A'-B'}] \succsim [x', A'; y', B'; g_{-A'-B'}]$.

Conditions requiring consistency of trade-offs across choices have been used before in the axiomatization of SEU; see in particular Wakker (1989). Our condition is actually more closely related to Ramsey's (1931) seminal contribution in which "tradeoffs"/"utility differences" are defined in terms of preferences over acts with two equally likely outcomes; indeed, Ramsey's axiom 2 is simply the restriction of Tradeoff Consistency to comparisons based on events of the form $B = A^c$ and $B' = A'^c$, that is: to events with unambiguous probability $\frac{1}{2}$.

4.2.2. The Representation Theorem.—

To endow Tradeoff Consistency with all its potential force, we shall also assume[28]

Axiom 17 (Solvability)

For any $x, y \in X$ *and* $T \in \Lambda^0$, *there exists* $z \in X$ *such that* $z \sim [x, T; y, T^c]$.

For expositional simplicity, assume also that the preference relation is **bounded** (in utility), i.e. that there exist $\overline{x}$ and $\overline{\overline{x}} \in X$ such that, for any $x \in X$, $\overline{x} \precsim x \precsim \overline{\overline{x}}$. Finally, we shall need to generalize Archimedean axiom to general acts.

Axiom 18 (Archimedean) *For any* $f, g \in \mathcal{F}$ *such that* $f \succ g$, *there exists* $K < \infty$ *and* $\frac{1}{K}$*−events* $C \subseteq \{\theta \mid f(\theta) \succ \overline{x}\}$ *and* $D \subseteq \{\theta \mid g(\theta) \prec \overline{\overline{x}}\}$, *such that, for any* $\frac{1}{K}$*−events*[29] C, D, $[f_{-C}, C^c; \overline{x}, C] \succ g$ *and* $f \succ [g_{-D}, D^c; \overline{\overline{x}}, D]$.

For consequences x, y such that $x \succ y$, let $\geq_{\{x,y\}}$ be the confidence relation associated with the DM's preferences over bets with these outcomes, and let $\overline{\rho} : \Sigma \rightarrow [0, 1]$ denote the confidence measure representing $\geq_{\{\overline{x}, \overline{\overline{x}}\}}$, with $\widehat{\overline{\rho}}$ as the canonical extension of $\overline{\rho}$ to $B(\Sigma, [0, 1])$. Note that $\widehat{\overline{\rho}}$ can be viewed as an "expectation operator" associated with the non-additive measure $\overline{\rho}$; for example, as remarked in section 3, $\widehat{\overline{\rho}}$ has the averaging property that $\widehat{\overline{\rho}}(c\mathbf{1}) = c$ for $c \in [0, 1]$.

[28] It is well-known how to derive this axiom from continuity plus connectedness assumptions.

[29] With respect to $\trianglerighteq^0$.

Theorem 4 *Consider a bounded preference ordering $\succsim$ on $\mathcal{F}$ that is compatible with an imprecise qualitative probability $\trianglerighteq^0$, and let $\overline{\rho}$ denote the representation of $\geq_{\{\overline{x},\overline{\overline{x}}\}}$. The preference ordering satisfies Tradeoff Consistency, the Archimedean axiom and Solvability if and only if there exists a (unique) u function u from X onto $[0,1]$ such that, for all $f,g \in \mathcal{F}$:*

$$f \succsim g \text{ iff } \widehat{\overline{\rho}}\,(u \circ f) \geq \widehat{\overline{\rho}}\,(u \circ g)\,. \tag{5}$$

In this case, if the acts f and g are salient unambiguous, then $f \succsim g$ iff $E_{\overline{\rho}} u \circ f \geq E_{\overline{\rho}} u \circ g$.

Preference relations described by Theorem 4 will be referred to as "**utility sophisticated**". By (5), utility sophisticated preferences over arbitrary acts are determined by the DM's preferences over bets and expected utility-preferences over salient unambiguous acts.

4.2.3. The Regular Case.—

U.s. preferences do not necessarily satisfy P4; they do, however, if and only if $\widehat{\overline{\rho}}$ is c-linear. This gives rise to the following Proposition.

Proposition 3 *A utility sophisticated preference ordering $\succsim$ satisfies P4 if and only if it*[30] *satisfies Union and Splitting Invariance.*

In view of Proposition 3, utility sophisticated preferences that satisfy P4 will be called "regular." In the regular case, $\overline{\rho}$ is represents the DM's confidence measure, and the bar can be dropped notationally. Since the operator $\widehat{\rho}$ is determined by the underlying ρ, it seems not entirely frivolous to write $\widehat{\rho}(Z)$ as $\int Z d\rho$.[31] We shall refer to this as the *intrinsic integral* of Z with respect to ρ, and sometimes write for clarity $\int_{int} Z d\rho$, to be distinguished from the Choquet integral $\int_{choq} Z d\rho$.

In general, the intrinsic integral is "complex" reflecting the complexity, lack of structure of the underlying confidence measure ρ. In the uniform model, however, $\widehat{\rho}$ has the following "interval expected utility" representation.

Fact 4 *If ρ is uniform with degree of caution γ, then*

$$\int_{int} u \circ f d\rho = \gamma \min_{\pi \in \Pi^*} E_\pi\,(u \circ f) + (1-\gamma) \max_{\pi \in \Pi^*} E_\pi\,(u \circ f)\,.$$

[30] More precisely, if all confidence relations $\geq_{\{x,y\}}$ satisfy these two conditions, or, equivalently, as is evident from the proof, if $\geq_{\{\overline{x},\overline{\overline{x}}\}}$ does.

[31] This notation cheats a bit in that it suppresses the reference relation $\trianglerighteq^0$ used to define $\widehat{\rho}$ from ρ. Mathematically, it would therefore seem desirable to show that the particular $\trianglerighteq^0$ used does not matter. If that can be done, it would furthermore be desirable to characterize compatibility of ρ with *some* minimally complete $\trianglerighteq^0$ internally.

In the general, non-uniform case, the intrinsic integral has a simple representation only for two-outcome acts, as noted in the following straightforward fact whose proof is omitted.[32] [33]

Fact 5 *If $f = [x, A; y, A^c]$ with $x \succsim y$, then*

$$\int_{int} u \circ f d\rho = u(x)\rho(A) + u(y)(1 - \rho(A)).$$

Under utility sophistication, the DM's revealed set of priors becomes highly informative for preferences over multi-valued acts. Indeed, Theorem 3, part 3), together with Theorem 4 entail the following characterization.

Proposition 4 *Suppose $\succsim$ is regular utility sophisticated. Then Π^* is the smallest set $\Pi \in \mathcal{K}(\Delta(\Sigma))$ such that*

$$f \succsim g \text{ whenever } E_\pi (u \circ f) \geq E_\pi (u \circ g) \text{ for all } \pi \in \Pi. \tag{6}$$

In particular, if the acts f and g are revealed unambiguous, then $f \succsim g$ iff $E_\rho u \circ f \geq E_\rho u \circ g$.

Thus Π^* is the smallest set of priors such that the DM's preferences coincide with the expected-utility preferences induced by the different priors, whenever these agree. Note that (6) is equivalent to requiring that there exists $\pi \in \Pi$ such that $E_\pi (u \circ f) > E_\pi (u \circ g)$ whenever $f \succ g$; thus, Π^* is the smallest set of priors Π that allows to rationalize each strict preference in terms of some prior chosen from Π.

In line with intuition, Proposition 4 has the following corollary.[34]

Corollary 1 *Regular utility sophistication implies Revealed Stochastic Dominance.*

Just as in the case of the regularity axioms Union and Splitting Invariance, the content of Tradeoff Consistency respectively utility sophistication should not depend on the particular imprecise qualitative probability $\trianglerighteq$ used in the range between $\trianglerighteq^0$ and $\trianglerighteq^*$ to be conceptually well-behaved. This follows directly from Proposition 4, from which it is evident that the preference relation is tradeoff-consistent with respect to $\trianglerighteq^*$. Thus we have[35]

[32] The proof is a straightforward consequence of the c-linearity of ρ.

[33] The class of preferences over two-outcome acts with this representation has been studied in detail in Ghirardato-Marinacci (2001b) under the name of "biseparability".

[34] This is an immediate consequence of Fact 3. Indeed, one obtains the stronger result that f revealed stochastichally dominates g if and only if there exists a u.s. preference relation $\succsim'$ whose restriction to bets and to comparisons of constant acts agrees with $\succsim$. The Corollary could also have been shown directly, without recourse to Propostion 4.

[35] Using Fact 2, one can also show that Tradeoff Consistency with respect to $\trianglerighteq^0$ together with Revealed Stochastic Dominance imply Tradeoff Consistency with respect to $\trianglerighteq^*$.

Corollary 2 *Regular Utilitarian Sophistication with respect to* $\trianglerighteq^0$ *implies Regular Utilitarian Sophistication with respect to* $\trianglerighteq^*$ *.*

4.2.4. Incompatibility of Utilitarian Sophistication with Choquet Expected Utility.— Utility sophistication is incompatible with the widely used Choquet Expected Utility (CEU) model: CEU preferences are utility sophisticated only if the DM maximizes subjective expected utility. We will state this fact formally as a statement about the relation of the Choquet and intrinsic integrals. Note that in general, the domain of applicability of the Choquet integral is larger than that of the intrinsic integral, in that the Choquet integral is defined for monotone transforms of the normalized confidence measure $\upsilon = \phi \circ \rho$, with ϕ reflecting the DM's probabilistic risk-attitude. However, ϕ must be the identity function if unambiguous acts are ranked according to their expected utility (as presupposed by utility sophistication).

Fact 6 $\int_{int} Zd\rho = \int_{choq} Zd\rho$ *if and only if* ρ *is a probability measure.*

In view of the conditional linearity property of the intrinsic integral stated as Lemma 2 in the appendix, mathematically Fact 6 is merely a restatement of the incompatibility of the one-stage and two-stage formulations of Choquet Expected Utility noted before in Sarin-Wakker (1992). Here, the incompatibility of the two methods of integration is not really that surprising since the underlying Cumulative Dominance respectively Tradeoff Consistency axioms are clearly very different ways of achieving the same thing, namely a determination of preferences over multi-valued acts through preferences over unambiguous acts and preferences over bets.

In view of the popularity and apparent usefulness of the CEU and the related Cumulative Prospect models in descriptive applications, the incompatibility may be viewed as disconcerting, and perhaps even as reason to question the appeal of Tradeoff Consistency. We see two potential interpretations of this incompatibility. On the one hand, the CEU model may be deemed to be an intrinsically *non-expected utility* model that is the simply not of interest under an "idealization" in which the DM is supposed to maximize expected utility over unambiguous acts. This interpretation does not seem very attractive, however, since there is nothing in the CEU model itself which precludes maximization of expected utility over unambiguous acts. A more convincing interpretation seems to be that the CEU model is appropriate only for subspaces on which the DM's beliefs have a rather specific structure, but that it is inappropriate if applied globally to preferences that are compatible with minimally complete probabilistic information[36]. This interpretation seems to be in accord with the apparent usefulness of the CEU model in descriptive applications, that is: in the stylized description of the outcome of laboratory experiments, since the state spaces in such experiments are typically extremely simple. Product- or conditional structures are rarely considered.

5. IS A FULLY BEHAVIORAL ACCOUNT POSSIBLE ?

5.1 Possibility in the Case of Utility Sophistication: Quantifying Out Salient Beliefs

Both the notion of revealed unambiguous beliefs as well as the axiom of Tradeoff Consistency underlying utility sophistication have been defined relative to an independently given imprecise qualitative probability $\trianglerighteq^0$. This relation has been interpreted as a set of likelihood judgements that is attributed to the DM directly rather than being inferred from his preferences alone; note however that, by compatibility of preferences with salient beliefs, the belief attribution must fail to be falsified by observable behavior. By contrast, a "fully behavioral" point of view would abstain from specifying a salient imprecise qualitative probability $\trianglerighteq^0$ as an independent primitive, and requires that all concepts be definable in terms of conditions on preferences alone. Is a fully behavioral account indeed possible ?

From a "fully behavioral" point of view, one can still make use of conditions such as Tradeoff Consistency by "quantifying out" the DM's salient belief relation $\trianglerighteq^0$. Specifically, one can determine from preferences whether there *exists* some minimally complete imprecise qualitative probability $\trianglerighteq^0$ relative to which preferences are utility sophisticated ; call such preferences *potentially utility sophisticated* . Prima facie, potential utility sophistication may be a weak and relatively unsatisfactory concept to the extent that there is little reason to attribute one of the associated imprecise qualitative probabilities to the DM. In particular, suppose that a given preference relation is compatible with both $\trianglerighteq_1^0$ and $\trianglerighteq_2^0$, but not with their union. Then, since there can be no evidence on preference grounds for privileging one over the other, there is little support for attributing *either* imprecise qualitative probability since there is no support for attributing them jointly. Fortunately, to due the specific content of utility sophistication, this issue in fact never arises. This follows from the following Proposition, which in turn follows directly from Proposition 4 and Corollary 2 above.

---

[36]Similar arguments have in fact been made before in Klibanoff (2001a,b) and Nehring (1999).

Proposition 5 *Suppose that there exist an imprecise qualitative probability* $\trianglerighteq^0$ *such that* $\succsim$ *is regularly utility sophisticated with respect to* $\trianglerighteq^0$ *. Then* $\succsim$ *is regularly utility sophisticated with respect to the revealed unambiguous belief relation* $\trianglerighteq^*$ *associated with* $\trianglerighteq^0$ *. Moreover,* $\trianglerighteq^*$ *contains any minimally complete imprecise qualitative probability* $\trianglerighteq^0$ *with respect to which* $\succsim$ *is regularly utility sophisticated .*

Thus, by Proposition 5, potentially utility sophisticated preferences are utility sophisticated relative to a unique maximal imprecise qualitative probability; label this imprecise qualitative probability (specified in the Proposition) $\trianglerighteq^{*_{us}}$. We suggest $\trianglerighteq^{*_{us}}$ as the canonical fully behavioral definition of revealed unambiguous beliefs. The superscript "$*_{us}$" indicates that this definition is applicable only to potentially utility sophisticated preferences; moreover, note that the imprecise qualitative probability $\trianglerighteq^{*_{us}}$ is determined from information about multi-valued preferences, not from preferences about bets alone. This is the price paid for achieving a fully behavioral definition of revealed unambiguous beliefs.[37]

5.2. Impossibility Without Utility Sophistication

Theorem 3 asserts for a regular, minimally unambiguous confidence relation $\geq$ the existence of a unique maximal imprecise qualitative probability $\trianglerighteq$ such that $\geq$ is compatible with it, and such that $\trianglerighteq$ extends $\trianglerighteq^0$, that is: such that revealed unambiguous beliefs are compatible with salient unambiguous beliefs. Theorem 3 does not exclude the possibility that there exists some other imprecise qualitative probability $\trianglerighteq'$ such that $\geq$ is compatible with it, while it is not compatible with any imprecise qualitative probability containing both $\trianglerighteq$ and $\trianglerighteq'$. Say that the confidence relation $\geq$ is *well-behaved* if no such $\trianglerighteq'$, i.e. if $\trianglerighteq^*$ is the unique maximal imprecise qualitative probability $\trianglerighteq$ such that $\geq$ is compatible with it. Label this imprecise qualitative probability $\trianglerighteq^{**}$, to denote its independence from the salient imprecise qualitative probability $\trianglerighteq^0$. For well-behaved confidence relations $\geq$, $\trianglerighteq^{**}$ suggests itself as a fully behavioral definition of revealed unambiguous beliefs, generalizing the definition of $\trianglerighteq^{*_{us}}$ in a natural way. Two questions arise naturally, a mathematical and a conceptual one.

The mathematical one asks *when* minimally unambiguous confidence relations are well-behaved (i.e. when $\trianglerighteq^{**}$ exists). The answer to this we do not know at this point. It seems likely that

[37] The possibility of defining a "revealed unambiguous preference" relation in fully behavioral terms partly motivated Ghirardato et al. (2001c) utility sophisticated point of view.

any uniform confidence relation is well-behaved; for all we know, it might even be that any regular minimally unambiguous confidence relation is well-behaved. (Note that in this case, Proposition 5 above would become redundant).

The second question is conceptual: If a minimally unambiguous confidence relation is well-behaved, does the relation $\trianglerighteq^{**}$ indeed provide a satisfactory *fully behavioral* definition of revealed unambiguous beliefs? This conceptual question becomes especially transparent in the case in which $\trianglerighteq^{**}$ is complete, i.e. when $\geq$ is a qualitative probability; under Revealed Stochastic Dominance (applied to $\trianglerighteq^{**}$), this is the case of "probabilistically sophisticated" preferences discussed in particular by Epstein (2001) and Ghirardato-Marinacci (2001a) . For specificity, consider preferences that have a CEU representation with Choquet capacity $\nu = \phi \circ \rho$, where ρ is a convex-ranged probability measure on Σ, and ϕ is a monotone, strictly convex mapping from $[0,1]$ onto $[0,1]$. Such preferences have two natural pure interpretations: on the one hand, the DM's belief may be free of ambiguity, being given by the additive subjective probability ρ, and the monotone transform ϕ may reflect his probabilistic risk-aversion. Alternatively, the DM could be probabilistically risk-neutral[38] but ambiguity averse and evaluate acts according to the minimum expected utility of the core of the capacity ν. In the first case, all events would be unambiguous, the second none. *On the basis of preferences alone, there seems to be no basis for preferring one interpretation; a satisfactory fully behavioral definition of unambiguous beliefs seems impossible.* At best, one can give a behavioral definition on the basis of a "convention", for example by declaring probabilistic sophistication to reveal absence of ambiguity *by definition.*[39]

Note that this second, conceptual question could be raised in principle also about $\trianglerighteq^{*us}$. But here, it lacks force. For example, if $\trianglerighteq^{*us}$ is complete as above, not only must $\geq$ be a subjective probability, but the DM must in fact be a SEU-maximizer; in this case, there is simply no basis for doubting the unambiguity of his beliefs. The case of utility sophistication is fundamentally simpler, since there is no need to disentangle ambiguity attitudes from probabilistic risk attitudes.

[38]And indeed utility sophisticated on a larger state-space in the manner of section 7.

[39]This in fact seems to be the line taken by Epstein-Zhang (2001). Ghirardato-Marinacci (2001a, section 6), on the other hand, argue for the convention aligned with the second interpretation, explaining very clearly the unavoidability of a conventional element in the absence of salient unambiguous events. They also point out how the existence of salient unambiguous events allows to distinguish between the two possible interpretations of a probabilistically sophisticated DM.

6. ELLSBERGIAN AMBIGUITY AVERSION

A central concern of the literature on decision making under ambiguity is the definition of an appropriate concept of "ambiguity aversion". The first and most frequently used is Schmeidler (1989)'s; it is motivated by the requirement that "substituting objective mixing for subjective mixing makes the decision maker better off" (p. 582).[40] In the Anscombe-Aumann-framework in which it was formulated, this amounts simply to convexity of preferences; for recent translations into a Savage framework, see Casadesus-Masanell et al. (2000) and Ghirardato et al. (2001d). Recently, Epstein (1999) as well as Ghirardato-Marinacci (2001a) have argued that the Schmeidlerian definition is intuitively unsatisfactory, and have proposed alternative definitions; both proposals derive a notion of "absolute" ambiguity aversion from a definition in terms of "comparative ambiguity aversion".

None of these definitions establishes a clear link between the formal definitions and the intuitions derived from the classical Ellsberg-type urn problems which presumably were the origin of the notion of ambiguity aversion in the first place. Ghirardato-Marinacci (2001a) establish the appropriateness of their definition to the Ellsberg 3-color problem; Epstein (1999, p.592) likewise uses that problem to illustrate the intuitive content of his definition but suspects that "a general formal result seems unachievable". In this section, we shall show that this lacuna can be remedied in the presence of minimally complete salient beliefs. By considering two classes of Ellsberg-style experiments appropriately generalized, we obtain natural analogues to the two types of definitions described above.

Ellsberg's two-urn problem motivates the following "minimal" definition of ambiguity aversion.[41]

Definition 3 $\geq$ *is* **minimally ambiguity averse** *if, for no* $A \in \Sigma$ *and* $T \in \Lambda^*$, $A > T$ *and* $A^c > T^c$.

Minimal ambiguity aversion precludes a simultaneous preference for betting on ambiguous event A over some unambiguous event T, combined with a preference for betting on the non-occurrence of A over the non-occurrence of T; this would contradict the typical outcome of Ellsberg's (1961) two-urn experiment, where exactly the opposite happens: the DM prefers to bet over any unambiguous event (T or T^c) over betting on any of the ambiguous event A and A^c.

Minimal ambiguity aversion is easily seen to be equivalent to requiring of the confidence measure that for all $A \in \Sigma : \rho(A) + \rho(A^c) \leq 1$. This definition is rather weak and does not appear very

[40] An adaptation of Schmeidler's definition to the present context is stated below as Axiom 21.

[41] In this and the following definitions, we rely on the revealed belief relation $\trianglerighteq^*$; in view of Fact 2 above, we could have used the DM's salient beliefs $\trianglerighteq^0$ instead throughout.

useful. Hence, in the following, we will consider more general definitions of ambiguity aversion that maintain the two basic features of the above distinctions: ambiguity aversion is defined in terms of betting preferences only, and more specifically, it imposes conditions on preferences comparisons between bets on ambiguous and unambiguous events.

A natural move is to consider urns with arbitrarily many colors.

Definition 4 *$\geq$ is* **weakly ambiguity averse** *if there do not exist partitions of Θ $\{A_i\}_{i=1,..,n}$ and $\{T_i\}_{i=1,..,n}$ such that $T_i \in \Lambda^*$ for all i, and such that $A_i > T_i$ for all $i \leq n$.*

In the language of urns, weak ambiguity aversion precludes the following preferences in an experiment in which a ball is drawn from each of two urns containing balls with n different colors labeled $i = 1, .., n$, a T-urn which contains the n colors in known frequencies, and an A-urn whose composition is "unknown". Then weak ambiguity aversion rules out that the DM, when asked to decide whether to bet on some color i as the outcome of a draw from urn A or to bet on the same color as that of the ball drawn from urn T, he prefers to bet on the ambiguous urn A over the unambiguous urn T, *for every color i*. Weak ambiguity aversion has the following attractive characterization.

Proposition 6 *For a regular confidence relation $\geq$, the following three conditions are equivalent.*

1. *$\geq$ is weakly ambiguity averse;*

2. *For any partition $\{A_i\}_{i=1,..,n}$ of Θ, $\sum_i \rho(A_i) \leq 1$;*

3. *there exists an additive probability measure $\pi \in \Delta(\Theta)$ such that $\rho(A) \leq \pi(A)$ for all $A \in \Sigma$.*[42]

In other words, weak ambiguity aversion is equivalent to ρ possessing a non-empty core. The key to proving the assertion is to note that under regularity, the second condition is equivalent to "balancedness" of ρ in the mixture-space extension, which is the classical characterizing condition of non-emptiness of the core; see, for example, Kannai (1992). Minimal unambiguity of preferences is critical to the validity of the characterization. Applied to preferences over bets, both the Epstein (1999) and Ghirardato-Marinacci (2001a) proposals amount to saying that a confidence relation is ambiguity averse if it is "more ambiguity averse" (in their sense) than some *precise* qualitative

[42]The proofs of the results in this section are fairly straightforward and will be supplied in future versions of this paper.

probability, which is equivalent to non-emptiness of the core.[43] Proposition 6 strengthens the appeal of their definition by showing that it captures a unified set of intuitions about ambiguity aversion. Moreover, weak ambiguity aversion can be tested directly, while a comparative definition cannot .

Yet weak ambiguity aversion does not *exhaust* our relevant intuitions, since it is limited a priori in its power by comparing ambiguous events to unambiguous ones only. Potential implications of "ambiguity aversion" (in a preformal sense) for comparisons of ambiguous events do not come into view, while they are essential to Schmeidlerian definitions of ambiguity aversion.

A more general viewpoint with an attendant stronger definition of ambiguity aversion is obtained by considering comparisons of ambiguous with conditionally unambiguous events. To motivate the following definition, consider an Ellsberg-style experiment with one urn containing five colors white, yellow, red, green, and black. The DM is told that there are 50 white or yellow, and 50 red or green balls, i.e. in total 100 non-black balls; he is not given any other information about the composition of the urn; in particular, he is not informed of the number of black balls. If given a choice between bets on any two non-black colors, an ambiguity-averse DM might reason as follows: my decision matters only if the ball drawn is not black. In that case, I have a fifty-fifty chance of winning if I bet on "white or yellow" (or on "red or green"), but only an ambiguous prospect otherwise. Hence the former bets are more attractive than the latter. In this example, it is thus natural to view the DM as displaying "ambiguity seeking" if instead he prefers to bet on "white or red" over betting on "white or yellow", *and* simultaneously prefers to bet on "yellow or green" over betting on "red or green".

This intuition is captured by the following general definition.

Definition 5 $\geq$ *is* **ambiguity averse** *if, for no* $A, A', T, T' \in \Sigma$ *such that* $A + A' = T + T'$ *and* $T \equiv^* T'$: $A > T$ *and* $A' > T'$.

One obtains the following representation theorem.

[43] In the CEU model with capacity $v = \phi \circ \rho$, hence with $\rho = \phi^{-1} \circ v$, this is just Epstein's characterization (Lemma 3.4, p. 589). Since the Ghirardato-Marinacci definition appeals an expected-utility benchmark, it is applicable here only under EU maximization over unambiguous acts, in which case $\phi = id$ and thus $v = \rho$; hence their characterization in terms of non-emptiness of the core of v is equivalent to that of the core of ρ.

Theorem 5 *For a regular confidence relation $\geq$, the following three conditions are equivalent.*

1. *$\geq$ is ambiguity averse;*
2. *ρ is super-additive, i.e. $\rho(A+B) \geq \rho(A) + \rho(B)$ for all A, B.*
3. *There exists $\Pi \in \mathcal{K}(\Delta(\Sigma))$ such that, for all $A \in \Sigma$,*

$$\rho(A) = \min_{\pi \in \Pi} \pi(A).$$

In this case, indeed $\Pi = \Pi^*$, and $\pi \in \Pi^*$ iff $\pi(A) \geq \rho(A)$ for all $A \in \Sigma$.

According to the first part of the Theorem, the "confidence" of an ambiguity-averse DM in an event is determined by its lower probability, relative to some belief set Π; according to the second part, this set must consist exactly of those probability measures that are compatible with his unambiguous beliefs. The first part is shown by noting that ambiguity aversion entails a weak form of preference convexity in the mixture-space extension, and by then appealing to Gilboa-Schmeidler's (1989) classical result. The second part is new; note that it implies that the revealed set of priors Π^* is uniquely identified by the associated lower-probability function $A \longmapsto \min_{\pi \in \Pi^*} \pi(A)$ it induces; this hinges critically on the range convexity of Π^*. [44]

In view of the subadditivity characterization in part 2) of Theorem 5, it is easily verified that under regularity, ambiguity aversion is equivalent to either of the following two conditions.

Axiom 19 (Complementarity)

For all $A, B \in \Sigma$ and $T, T' \in \Lambda^*$: $A \doteqdot T$ and $B \doteqdot T'$ imply $A + B \geq T + T'$.

This condition captures the intuition that the ambiguities of disjoint events can never reinforce each other, but that they can cancel each other out.

Axiom 20 (Preference for Randomization over Bets)

For any $A, B \in \Sigma$ such that $A \geq B$ and $T \in \Lambda^*$ such that $T \cap D \equiv^* T^c \cap D$ for any $D \in \{A \backslash B, A+B, B \backslash A, (A+B)^c\}$: $(T \cap A) + (T^c \cap B) \geq B$.

Here the event T is specified to have conditional probability $\frac{1}{2}$ irrespective of the joint realization of A and B; thus the event $(T \cap A) + (T^c \cap B)$ can be viewed as describing a random bet that is paid out in the event A or in the event B, contingent on the outcome of the "fair coin toss" T.

[44] In a finite state setting, for example, only very special classes of beliefs are generated by their lower-probability function; see, for example, Walley (1991, section 4.6.1)

Preference for Randomization over Bets can be viewed as a version of ambiguity aversion a la Schmeidler restricted to preferences over bets. Schmeidler's original definition which applies to general multi-valued acts it can be reformulated here as follows.

Axiom 21 (Preference for Randomization over Multi-Valued Acts)

For any $f, g \in \mathcal{F}$ such that $f \succsim g$ and any $T \in \Lambda^*$ such that $T \cap D \equiv^* T^c \cap D$ for all D contained in the algebra generated by f and g : $[f, T; g, T^c] \succsim g$.

If all departures from SEU maximization are due to ambiguity, i.e. if preferences are "utility sophisticated", then our notion of ambiguity aversion entails Schmeidler's; this follows from Proposition 7 below.[45] Otherwise, it is substantially weaker. For example, our definition is entirely consistent with the CEU model, while Schmeidler's is not, as shown in Klibanoff (2001a).

Note also that ρ viewed as a capacity (normalized to be additive on salient unambiguous events Λ^0 as assumed throughout) need not be "convex" but merely "exact", that is: to be a lower probability in the sense of Theorem 5. Indeed, one can show that if ρ is convex and regular with respect to a minimally complete $\trianglerighteq^0$, then it must be additive. Hence the misgivings of both Ghirardato-Marinacci (2001a) and Epstein (1999) about the counterintuive convexity implications on the capacity entailed by a Schmeidlerian definition do not apply to the present one, while their misgivings about convexity itself are fully borne out.

From Theorems 5 and 4, we obtain the following characterization of the classical Minimum Expected Utility (MEU) model which is given by the following representation:

$$f \succsim g \text{ if and only if } \min_{\pi \in \Pi} E_\pi (u \circ f) \geq \min_{\pi \in \Pi} E_\pi (u \circ g),$$

for appropriate utility functions u and belief sets Π.

Proposition 7 *Let $\succsim$ be a solvable, minimally unambiguous preference ordering. Then $\succsim$ has a Minimum Expected Utility representation if and only if it is ambiguity averse, tradeoff consistent, Archimedean and satisfies P4.*

The MEU model has been axiomatized first in the Anscombe-Aumann-framework by Gilboa-Schmeidler (1989), and recently in a Savage framework by Casadesus-Masanell et al. (2000) and

[45] Klibanoff (2001b), for example, is explicit about the implicit "utility sophisticated" character of Schmeidler's notion by saying that "one may interpret this requirement as saying that the individual likes smoothing expected utility across states" (p. 290).

Ghirardato et al. (2001d). Proposition 7 has advantages in the conceptual transparency and minimality of its behavioral/rationality assumptions, especially in appealing to P4 rather than "certainty-independence" type assumptions, and by relying on a notion of ambiguity aversion defined in terms of betting preferences; on the other hand, it pays a price in the form of stronger structural assumptions: Gilboa-Schmeidler (1989) do not need solvability, while Casadesus-Masanell et al. (2000) and Ghirardato et al. (2001d) make do without a richness of states requirement.

7. SMALL STATE SPACES

7.1. Motivation: "Descriptive" Applications

In the literature, richness assumptions are frequently criticized.[46] This criticism is viewed to have particular force in "descriptive" applications; see, for example Epstein (1999). While the meaning of "descriptive" is often not entirely clear[47], the underlying intention seems to be that typically, only preferences over acts defined in terms of a limited set of states are "observed" or specified in an economic model. Hence it is deemed desirable to define and analyze fundamental concepts such as ambiguity aversion or unambiguous beliefs / events in terms of "observed" behavior/preferences in the context of the given "observed" state space. That state space may be small; for example, in the analysis of Ellsberg experiments, the relevant state space may consist merely of the possible colors of the balls in the urn.[48]

While we concur with the interest of exploring the behavioral *consequences* of fundamental notions in specific, possibly very restricted "observed contexts", it does not seem to be appropriate to confine oneself from the outset to such contexts if the task is to clarify fundamental conceptual questions, especially if they have normative content. To address them, "richness" of the framework is generally desirable, as it gives maximum play to the substantive assumptions. Indeed, richness assumption may prove to be indispensable; for example, as argued in section 5, it simply does not seem to be possible to provide a sound general definition of "revealed unambiguous beliefs" without them.

For conceptual and normative purposes, it is entirely sufficient that preferences be "observable

[46]With respect to the AA approach in particular, these express unhappiness about the inclusion of "extraneous" random devices.

[47]For Epstein (1999, p. 601), for example, observability in laboratory experiments does not constitute "descriptiveness" per se.

[48]Note that if actual observ*ed*-ness is taken seriously, it seems clear in any case that the state space needs to be finite. Hence, according to this line of criticism, even Savage's own theory goes beyond what is actually "observed".

in principle", that is: correspond to choices in well-defined *hypothetical* choice situations; actual observedness is not needed. As a result, a richness assumption such as Equivisibility seems perfectly viable; for example, fleshing it out in Anscombe-Aumann manner in terms of an extraneous continuous random device is a way of constructing such well-defined hypothetical choice situations.[49]

7.2. In the Background: A Rich State-Space

To integrate the two viewpoints, the "conceptual" and the "descriptive", we shall now assume that the DM "has" a minimally unambiguous preference relation that is observable in principle, but that the analyst *observes* only preferences over acts measurable with respect to some coarser algebra Σ_{obs}. For specificity, assume Σ_{obs} to be the algebra generated by some finite partition $\mathcal{S}$ of Θ; the elements of the partitions correspond to the observable "states". The observable preference relation $\succsim_{obs}$ is the restriction of $\succsim$ to $\mathcal{F}(\Sigma_{obs}) \times \mathcal{F}(\Sigma_{obs})$. Assume that the DM's underlying preferences $\succsim$ are regular utility sophisticated . In view of Theorem 4 and Proposition 3, his observable preferences $\succsim_{obs}$ have a representation of the form

$$f \succsim_{obs} g \text{ if and only if } I(u \circ f) \geq I(u \circ g), \tag{7}$$

where $u : X \to \mathbf{R}$ has range $[0,1]$ and $I : [0,1]^{\mathcal{S}} \to \mathbf{R}$ is monotone and c-linear. Indeed, I is simply the restriction of $\widehat{\rho}$ to $B(\Sigma_{obs}, [0,1])$.

While it will not be possible in general to recover the entire underlying preference relation $\succsim$, it turns out to be possible to recover a substantial, *minimally unambiguous* fragment of that relation, namely its restriction to acts measurable with respect to the "mixture-space (of events) generated by Σ_{obs}" $\mathcal{M}(\Sigma_{obs})$. $\mathcal{M}(\Sigma_{obs})$ is defined as follows:

$$\mathcal{M}(\Sigma_{obs}) := \{A \in \Sigma \mid A \cap S \in \Lambda^0_S \text{ for all } S \in \mathcal{S}\}.$$

Let $\succsim_{\mathcal{M}(\Sigma_{obs})}$ and $\trianglerighteq^0_{\mathcal{M}(\Sigma_{obs})}$ denote the restrictions of $\succsim$ respectively $\trianglerighteq^0$ to $\mathcal{M}(\Sigma_{obs})$.

Note that by construction, $\mathcal{M}(\Sigma_{obs})$ is a λ-system, and that $\succsim_{\mathcal{M}(\Sigma_{obs})}$ is compatible with $\trianglerighteq^0_{\mathcal{M}(\Sigma_{obs})}$ which is minimally complete. To see that $\succsim_{\mathcal{M}(\Sigma_{obs})}$ can be uniquely inferred from $\succsim_{obs}$, one merely

[49] Indeed, this view seems to concord with the mainstream of the decision theoretic tradition following Ramsey, deFinetti, and Savage. In particular, the central notion of "revealed preference" in that tradition requires definition/construction of concepts in terms of choices that are *observable in principle* rather than *observed de facto* in any serious sense. (It goes without saying that this discussion deserves more detail in its own right, and thus must be taking with a grain of salt in its present form. Here, its purpose is essentially motivational, to provide a starting point and context for the subsequent analysis.)

needs to verify that the restriction of ρ to $\mathcal{M}(\Sigma_{obs})$ can be uniquely inferred by utility sophistication. But this follows from the identity

$$\rho(A) = \rho(\sum_{S\in\mathcal{S}} A\cap S) = \widehat{\rho}\left(\sum_{S\in\mathcal{S}} \pi_S^0 (A\cap S)\, 1_S\right) = I\left(\sum_{S\in\mathcal{S}} \pi_S^0 (A\cap S)\, 1_S\right).$$

In the Anscombe-Aumann example of section 2, if Σ_{obs} is the algebra Σ_1, then $\mathcal{M}(\Sigma_{obs}) = \Sigma_1 \times \Sigma_2$; thus, in this case, the entire preference relation can be recovered from the directly observed preferences.

The upshot of this discussion is that preferences over the rich state $\mathcal{M}(\Sigma_{obs})$ can be inferred from observable choice behavior *once it is known that the DM's full preferences are utility sophisticated.* As a result, the central analytic concepts remain well-defined *in terms of observed preferences* even when Σ_{obs} is small.

7.3. Priors Revealed on Small State Spaces

We shall now apply these ideas to characterizing unambiguous beliefs in terms of observed preferences under the assumption of utility sophistication. Since with small state spaces it is generally not possible to characterize beliefs in terms of an imprecise qualitative probability, one must do this directly in terms of a multi-prior representation.

Since the DM's beliefs on the rich space $\mathcal{M}(\Sigma_{obs})$ can be inferred from observable preferences, it is natural to define the "observable set of priors" $\Pi^{us}_{\succsim_{obs}}$ as the restriction of the set of priors (revealed on $\mathcal{M}(\Sigma_{obs})$) $\Pi_{\left(\geq_{\mathcal{M}(\Sigma_{obs})}\right)^*} \in \mathcal{K}(\Delta(\mathcal{M}(\Sigma_{obs})))$ to the observable event space Σ_{obs}. To define it formally, let $rest_{\Sigma'}\Pi$ as $\{rest_{\Sigma'}\pi \mid \pi\in\Pi\}$, where $rest_{\Sigma'}\pi$ is the restriction of π to $\Sigma' \subseteq \Sigma$.

Definition 6 $\Pi^{us}_{\succsim_{obs}} := rest_{\Sigma_{obs}}\Pi_{\left(\geq_{\mathcal{M}(\Sigma_{obs})}\right)^*}$.

In view of Proposition 4, the observable set of priors $\Pi^{us}_{\succsim_{obs}}$ is the smallest set of priors such that the DM's preferences coincide with the expected-utility preferences induced by the different priors, whenever these agree. This is stated by the following Proposition.

Proposition 8 $\Pi^{us}_{\succsim_{obs}}$ *is the smallest set* $\Pi \in \mathcal{K}(\Delta(\Sigma_{obs}))$ *such that*

$$f \succsim_{obs} g \text{ whenever } E_\pi f \geq E_\pi g \text{ for all } \pi \in \Pi. \tag{8}$$

Conceptually, the set $\Pi^{us}_{\succsim_{obs}}$ describes the DM's set of priors over the observed state space Σ_{obs} inferred from his entire observed preference relation on the basis of his assumed utility sophistication.

Note that by Proposition 8 $\Pi^{us}_{\succsim_{obs}}$ does not depend on the salient belief relation $\trianglerighteq^{0}_{\mathcal{M}(\Sigma_{obs})}$ implicitly used in its definition.

7.4. From Unambiguous Beliefs to Unambiguous Preferences.

Proposition 8 motivates the following definition of the DM's "unambiguous preference" subrelation $\succsim^*_{obs}$ as those preferences that are consistent with expected utility maximization relative to $\Pi^{us}_{\succsim_{obs}}$.

Definition 7 *$f \succsim^*_{obs} g$ iff $E_\pi (u \circ f) \geq E_\pi (u \circ g)$ for all $\pi \in \Pi^{us}_{\succsim_{obs}}$.*

The unambiguous preference subrelation $\succsim^*_{obs}$ (hence $\Pi^{us}_{\succsim_{obs}}$ as well) can be characterized in behavioral terms directly, paralleling Theorem 3. To do this, assume for simplicity that the DM is risk-neutral, i.e. that $X = [0,1]$ and $u = id$. As in section 3, $\succsim'$ is **independent** if $f \succsim' g$ if and only if $\alpha f + (1-\alpha)h \succsim' \alpha g + (1-\alpha)h$ for all f, g, h and $\alpha \in (0,1]$.[50]

Proposition 9 *Suppose that $\succsim_{obs}$ is c-linear, continuous and risk-neutral. Then $\succsim^*_{obs}$ is the maximal independent subrelation of $\succsim_{obs}$. Moreover, $f \succsim^*_{obs} g$ if and only if $\alpha f + (1-\alpha)h \succsim_{obs} \alpha g + (1-\alpha)h$ for all h and $\alpha \in (0,1]$.*

The first characterization of $\succsim^*_{obs}$ as the maximal independent subrelation of $\succsim_{obs}$ has been proposed before as a *definition* of unambiguous preferences in an Anscombe-Aumann framework in Nehring (1996), which also contained a variant of the second characterization.[51]

The present approach improves upon these two contributions first of all in its greater generality, by providing a single unified definition of unambiguous beliefs which induces the above notion if preferences are utility sophisticated. It has significant advantages even if utility sophistication is assumed as a maintained (rationality) assumption. In particular, utility sophistication can be formulated explicitly and hence behaviorally verified only in a rich setting via Tradeoff Consistency. By contrast, in a small state-space[52], "utility sophistication" can be only appealed to as an informal, interpretative assumption to the effect that all departures from EU maximization are attributable to ambiguity. Furthermore, the relation $\succsim^*_{obs}$ itself is a *preference* relation in contrast to $\trianglerighteq^*$ which *is* a belief relation.

Paralleling Proposition 1, one can characterize unambiguous events in a variety of equivalent ways.

[50] In the general, non-risk-neutral case, independence can be defined in the manner of Ghirardato et al. (2001d).

[51] The exact form of the second characterization has also recently been arrived at independently by Ghirardato et al. (2001c).

[52] That is: in a small state space that is not extended to a rich one in the manner of subsection 7.2.

Proposition 10 *Suppose that* $\succsim_{obs}$ *is c-linear, continuous and risk-neutral. Then the following three statements are equivalent .*

1. $\pi(A) = \pi'(A)$ *for all* $\pi, \pi' \in \Pi^{us}_{\succsim_{obs}}$.
2. *For all constant acts* f, $1_A \succsim^*_{obs} f$ *or* $f \succsim^*_{obs} 1_A$.
3. *For all acts* f, g *and* f', g' *such that* $f - g$ *is* $\{A, A^c\}$*-measurable and* $f - g = f' - g'$, $f \succsim_{obs} g$ *if and only if* $f' \succsim_{obs} g'$.

The second condition has been independently proposed as a definition of revealed unambiguous events by Ghirardato et al. (2001c), who also show its equivalence to the first. The third characterization of an unambiguous event has been proposed before in Nehring (1999). It has the following intuitive interpretation. Think of $f - g$ as the *incremental bet* involved in deciding between f and g; if indeed the DM assigns an unambiguous probability to the event A, the incremental bet has an unambiguous expectation which determines the ranking between f and g, as well as the parallel ranking between f' and g'.

APPENDIX: PROOFS.

Proof of Fact 1.

Union Invariance equivalent to c-additivity. That c-additivity implies Union Invariance is straightforward. For the converse, write $Y = \sum_{i \in I} y_i 1_{E_i}$. Since $Y \leq (1-c)\mathbf{1}$, there exist $A \in [Y]$ and $S, T \in \Lambda^0$ such that $\rho(S) = \rho(A) \leq c$, $\rho(T) = c$, and T is disjoint from both A and S. To see this, take $A = \sum_{i \in I} A_i$ with $A_i \in \Lambda^0_{E_i}$ and $\pi^0(A_i/E_i) = y_i$, $S = \sum_{i \in I} S_i$ with $S_i \in \Lambda^0_{E_i}$ and $\pi^0(S_i/E_i) = \rho(A)$, and $T = \sum_{i \in I} T_i$ with $T_i \in \Lambda^0_{E_i}$ and $\pi^0(S_i/E_i) = c$ such that T_i is disjoint from both A_i and S_i, for all $i \in I$; such A_i, S_i, and T_i exist by the range convexity of Π^0. Clearly, $A + T \in [Y + c\mathbf{1}]$. Since $A \equiv S$ by assumption, $A + T \equiv S + T$ by Union Invariance; which is tantamount to $\rho(A+T) = \rho(S+T) = \rho(S) + \rho(T) = \rho(A) + c$. Hence

$$\widehat{\rho}(Y + c\mathbf{1}) = \rho(A+T) = \rho(A) + c = \widehat{\rho}(Y) + c.$$

Splitting Invariance equivalent to positive homogeneity.

It is clear that positive homogeneity implies Splitting Invariance; to show the converse, a similar technique as in part establishes that Splitting Invariance implies positive homogeneity for rational α. This implies positive homogeneity for arbitrary α, since by monotonicity of $\widehat{\rho}$,

$$\alpha\widehat{\rho}(Y) = \sup\{\widehat{\rho}(\beta Y) \mid \beta \leq \alpha, \beta \in \mathbf{Q}\} \leq \widehat{\rho}(\alpha Y) \leq \inf\{\widehat{\rho}(\beta Y) \mid \beta \geq \alpha, \beta \in \mathbf{Q}\} = \alpha\widehat{\rho}(Y),$$

and thus $\widehat{\rho}(\alpha Y) = \alpha\widehat{\rho}(Y)$. □

Proof of Theorem 3. Consider the extension of $\widehat{\geq}$ of $\geq$ to the mixture-space $B(\Sigma, [0,1])$ associated with $\trianglerighteq^0$. Define the following subrelation of $\widehat{\gtrsim}$:

$$X \widehat{\gtrsim}^* Y \text{ if and only if, for all } n \in \mathbf{N} \text{ and } Z \in B(\Sigma, [0,1]) : \frac{1}{n}X + \frac{n-1}{n}Z \widehat{\gtrsim} \frac{1}{n}Y + \frac{n-1}{n}Z.$$

Lemma 1 *If $X \widehat{\gtrsim}^* Y$, then $X' \widehat{\gtrsim}^* Y'$ for all X', Y' such that $X' - Y' = \frac{1}{\gamma}(X - Y)$ for some $\gamma \in \mathbf{R}_{++}$.*

Consider X, Y such that $X \widehat{\gtrsim}^* Y$. By continuity, it suffices to prove the assertion for rational γ.

To verify it, define

$$Z(\beta) = \frac{1}{1 - \beta\gamma}\left[\beta(X' - \gamma X) + (1-\beta)\frac{1}{2}\mathbf{1}\right].$$

Clearly, for sufficiently small β, $Z(\beta) \in B(\Sigma, [0,1])$.

Now define

$$X'' = \beta\gamma X + (1-\beta\gamma) Z(\beta), \text{ and} \quad (9)$$
$$Y'' = \beta\gamma Y + (1-\beta\gamma) Z(\beta).$$

With a bit of algebra, it is easily verified that

$$X'' = \beta X' + (1-\beta)\frac{1}{2}\mathbf{1}, \text{ and} \quad (10)$$
$$Y'' = \beta Y' + (1-\beta)\frac{1}{2}\mathbf{1}.$$

Take a sufficiently small β with the property that $\beta\gamma = \frac{1}{n}$ for appropriate n. By the assumption and (9), $X''\widehat{\geq}Y''$. Hence by c-independence and (10), $X'\widehat{\geq}Y'$, verifying the Lemma.

From the Lemma, it follows immediately that $\widehat{\geq}^*$ satisfies independence, hence that it is the maximal independent subrelation of $\geq$.

Note that one can write $\widehat{\geq}^*$ as $\cap\widehat{\geq}_{n,Z}$, where $\widehat{\geq}_{n,Z}$ is given by $X\widehat{\geq}_{n,Z}Y$ iff $\frac{1}{n}X+\frac{n-1}{n}Z\widehat{\geq}\frac{1}{n}Y+\frac{n-1}{n}Z$. Thus $\widehat{\geq}^*$ is transitive and continuous, since these properties are preserved by taking intersections; it is evidently monotone as well.

Translating $\widehat{\geq}^*$ back into Σ yields the ordering $\trianglerighteq^*$ described in part ii) of the Theorem. By the above results, it is the maximal imprecise qualitative probability compatible with $\geq$ and containing $\geq_0$.

As to part iii), it follows from Theorem 2 and part i) of this Theorem that Π^* is the smallest closed, convex set Π such that $A \geq B$, whenever $\pi(A) \geq \pi(B)$ for all $\pi \in \Pi$. By modus tollens and the completeness of $\geq$, the latter is however equivalent to the condition that if $A > B$, there exists $\pi \in \Pi$ such that $\pi(A) > \pi(B)$, as desired. $\square$

Proof of Proposition 1.

i) 1. is equivalent to 2.: straightforward.

ii) 1. is equivalent to 3. Suppose $A \equiv^* T$ for some $T \in \Lambda^0$. By the (proof of) Theorem 3, then also $A' + C \doteqdot T' + C$ for any $A' \in \Lambda^0_A$ and $T' \in \Lambda^0_T$ such that $\pi^0(A'/A) = \pi^0(T'/T) =: \alpha \in (0,1]$. Hence by c-linearity, $\rho(A'+C) = \rho(T'+C) = \rho(T') + \rho(C) = \rho(A') + \rho(C)$. The converse implication is obtained by reverting this inference, quantifying over A' and C.

iii) 3. is equivalent to 4.

It is straightforward that 3. implies 4. To verify the converse, suppose that 3. is violated, i.e. that there exists $A, B \in \Sigma, A' \in \Lambda^0_A$ disjoint from B, and $\varepsilon \neq \emptyset$ such that

$$\rho(A' + B) = \rho(A') + \rho(B) + 2\varepsilon. \tag{11}$$

Consider the case of $\varepsilon > 0$; the case of $\varepsilon < 0$ is similar. Pick $n \in \mathbf{N}$ such that $\frac{1}{n} < 1 - \rho(B) - \varepsilon$, and pick $A'' \in \Lambda^0_{A'}$ and $B'' \in \Lambda^0_{B'}$ such that $\pi^0(A''/A') = \pi^0(B''/B') = \frac{1}{n}$. Moreover, take $T \in \Lambda^0$ disjoint from A'' such that

$$\rho\left(T\right) = \frac{1}{n}\left(\rho(B) + \varepsilon\right); \tag{12}$$

such T exists by the range convexity of Π^0. By c-linearity,

$$\rho(A'' + T) = \frac{1}{n}\rho(A') + \rho\left(T\right). \tag{13}$$

Likewise, since $A'' + B'' \in \Lambda^0_{A'+B}$ with $\pi^0(A'' + B''/A' + B)$ by construction, by positive homogeneity and (11) one has

$$\rho(A'' + B'') = \frac{1}{n}\rho(A' + B) = \frac{1}{n}\left(\rho(A') + \rho(B) + 2\varepsilon\right). \tag{14}$$

By (12) and positive homogeneity, $\rho\left(T\right) > \frac{1}{n}\rho(B) = \rho(B'')$, and thus

$$T > B''.$$

On the other hand, by (13) and (14), $\rho(A'' + T) - \rho(A'' + B'') = \frac{1}{n}\rho(A') + \frac{1}{n}\left(\rho(B) + \varepsilon\right) - \frac{1}{n}\left(\rho(A') + \rho(B) + 2\varepsilon\right) = -\frac{1}{n}\varepsilon < 0$, hence

$$A'' + T > A'' + B'',$$

the desired violation of 4., since $A'' \in \Lambda^0_A$ as is easily verified. □

Proof of Theorem 4.

(Necessity)

That the representation (5) entails Tradeoff Consistency follows from the *Conditional Linearity* property of $\widehat{\rho}$ stated in the following Lemma. Say that $Z \in B(\Sigma, [0,1])$ is $\trianglerighteq^0$-unambiguous conditional on the finite partition $\mathcal{S}$ if, for all $S_i \in \mathcal{S}$, $Z 1_{S_i}$ is $\Lambda^0_{S_i}$-measurable; let $B^0(\Sigma/\mathcal{S}, [0,1])$ denote their class. For $Z \in B^0(\Sigma/\mathcal{S}, [0,1])$, the expectation conditional on $\mathcal{S}$ $E^0(Z/\mathcal{S})$ is a well-defined random variable given by

$$E^0(Z/\mathcal{S}) := \sum_{S_i \in \mathcal{S}} 1_{S_i} \left(\sum_{z \in [\mathbf{0},\mathbf{1}]} z\pi^0(\{\theta \in S_i \mid Z\left(\theta\right) = z\}/S_i) \right).$$

Lemma 2 *(Conditional Linearity) For $Z \in B^0(\Sigma/\mathcal{S}, [0,1])$, $\widehat{\rho}(Z) = \widehat{\rho}\left(E^0(Z/\mathcal{S})\right)$.*

Note that the Lemma asserts in particular that $\widehat{\rho}$ restricted to unambiguous random variables is the ordinary expectation with respect to π^0 or equivalently $\overline{\rho}$.

To verify Lemma 2, write Z as $\sum_{i,j} z_{ij} 1_{A_{ij}}$ with $S_i = \sum_{j \le n_j} A_{ij}$ for all i. Consider any $C \in [Z]$ such that $\pi(C \cap A_{ij}) = z_{ij}\pi(A_{ij})$ for all i, j and all $\pi \in \Pi$. Then in fact, for all i and all $\pi \in \Pi$, $\pi(C \cap S_i) = \sum_j \pi(C \cap A_{ij}) = \sum_j z_{ij}\left[\pi^0(A_{ij}/S_i)\,\pi(S_i)\right] =$ $\left[\sum_j z_{ij}\pi^0(A_{ij}/S_i)\right]$, which implies that $C \in [E^0(Z/\mathcal{S})]$. Thus indeed $C \in [Z] \cap [E^0(Z/\mathcal{S})]$, which establishes the Lemma in view of the assumed compatibility of $\overline{\rho}$ with $\trianglerighteq^0$. $\square$

Necessity of the other conditions is straightforward.

(Sufficiency).

By the Archimedean axiom, Solvability and compatibility with $\trianglerighteq^0$, it is straightforward to see that there exists a (unique) monotone functional $I : B(\Sigma, [0,1]) \to [0,1]$ and a (unique) function u from X onto $[0,1]$ such that

i) $I(\alpha\mathbf{1}) = \alpha$ for all $\alpha \in [0,1]$,

ii) $I(1_A) = \overline{\rho}(A)$ for all $A \in \Sigma$,

iii) $f \succsim g$ iff $I(u \circ f) \ge I(u \circ g)$ for all $f, g \in \mathcal{F}$.

By i), ii) and iii), for any $x \in X$, $u(x) = \overline{\rho}(T)$ for any T such that $x \sim [\overline{x}, T^c; \overline{\overline{x}}, T]$.

To show that in fact $I = \widehat{\rho}$, we shall first consider the case of dyadic-valued utilities; a number is *dyadic* if $\alpha = \frac{\ell}{2^m}$, where m is natural or zero, and ℓ is an odd integer or zero; m will be referred to as the (dyadic) order of α denoted by $|\alpha|$. Let $\mathbf{D}$ denote the set of dyadic numbers in $(0,1]$.

Lemma 3 *For any $\alpha \in \mathbf{D}, w, x, y \in X, A, B, T \in \Sigma$ such that $\pi^0(T) = \pi^0(A/B) = \alpha$: if $w \sim [x, T; y, T^c]$, then $[w, B; f_{-B}, B^c] \sim [x, A; y, B\backslash A; f_{-B}, B^c]$.*

The Lemma is proved by induction on the order of α. If the order of α is 1, i.e. if $\alpha = \frac{1}{2}$, the assertion follows directly from Tradeoff Consistency. Suppose thus, that the Lemma has been shown for all instances in which the order of the dyadic coefficient α' is strictly less than that of α. Assume that $\alpha \ge \frac{1}{2}$; the case of $\alpha < \frac{1}{2}$ can be proved essentially identically. Then $\alpha = \frac{1}{2} + \frac{1}{2}\beta$, where β is dyadic with $|\beta| = |\alpha| - 1$.

Now define events T_1, T_2, T_3 such that $T_1 + T_2 + T_3 = \Theta, T_2 + T_3 = T$, and $\pi^0(T_2) = \frac{1}{2}\beta$, hence also $\pi^0(T_3) = \frac{1}{2}$ and $\pi^0(T_2/T_1 + T_2) = \beta$.

Likewise, define events A_1, A_2, A_3 such that $A_1 + A_2 + A_3 = B, A_2 + A_3 = A$, and $\pi^0(A_2/B) = \frac{1}{2}\beta$, hence also $\pi^0(A_3/B) = \frac{1}{2}$ and $\pi^0(A_2/A_1 + A_2) = \beta$. Such events exist by the range convexity of $\Pi_{\trianglerighteq 0}$.

Take any $D \in \Lambda^0$ such that $\pi^0(D) = \beta$, and $z \in X$ such that $z \sim [x, D; y, D^c]$; such z exists by Solvability. By the induction assumption, $[z, T_1 + T_2; x, T_3] \sim [y, T_1; x, T_2; x, T_3]$, hence by the assumption on w also

$$[z, T_1 + T_2; x, T_3] \sim [w, T_1 + T_2; w, T_3]. \tag{15}$$

Writing $[x, A; y, B\backslash A; f_{-B}, B^c] = [y, A_1; x, A_2; x, A_3; f_{-B}, B^c]$, this act is indifferent to $[z, A_1 + A_2; x, A_3; f_{-B}, B^c]$ by induction assumption, which in turn is indifferent to $[w, A_1 + A_2; w, A_3; f_{-B}, B^c]$ by Tradeoff Consistency and (15). By transitivity, we get

$$[x, A; y, B\backslash A; f_{-B}, B^c] \sim [w, B; f_{-B}, B^c],$$

as desired. □

By the lemma, we obtain the desired conclusion for dyadic-valued functions $B(\Sigma, \mathbf{D} \cup \{0\})$, which we shall abbreviate to $B_{\mathbf{D}}$. Thus, take any $Y \in \mathbf{D} \cup \{0\}$; by solvability, there exists $f = [w_i, B_i]_{i \leq n} \in \mathcal{F}$ such that $Y = \sum_{i \leq n} u(w_i) 1_{B_i} = u \circ f$. For each $i \leq n$, pick $A_i \subseteq B_i$ such that $\pi^0(A_i/B_i) = u(w_i)$. By $n-$fold application of Lemma 3, $f \sim \left[\overline{\overline{x}}, \sum_{i \leq n} A_i; \overline{x}, \left(\sum_{i \leq n} A_i\right)^c\right]_{i \leq n}$. Hence $I(Y) = I(u \circ f) = \overline{\rho}(\sum_{i \leq n} A_i) = \widehat{\overline{\rho}}(Y)$; the latter follows since $\sum_{i \leq n} A_i \in [Y]$ by construction.

Thus $I = \widehat{\overline{\rho}}$ on $B_{\mathbf{D}}$. To extend this to all of $B(\Sigma, [0, 1])$, one can reason as follows.

By monotonicity,

$$\sup\{I(Z) \mid Z \in B_{\mathbf{D}}, Z \leq Y) \leq I(Y) \leq \inf\{I(Z) \mid Z \in B_{\mathbf{D}}, Z \geq Y),$$

hence by the validity of $I = \widehat{\overline{\rho}}$ on $B_{\mathbf{D}}$ also

$$\sup\{\widehat{\overline{\rho}}(Z) \mid Z \in B_{\mathbf{D}}, Z \leq Y) \leq I(Y) \leq \inf\{\widehat{\overline{\rho}}(Z) \mid Z \in B_{\mathbf{D}}, Z \geq Y).$$

By monotonicity and continuity of $\widehat{\overline{\rho}}$,

$$\sup\{\widehat{\overline{\rho}}(Z) \mid Z \in B_{\mathbf{D}}, Z \leq Y) = \widehat{\overline{\rho}}(Y) = \inf\{\widehat{\overline{\rho}}(Z) \mid Z \in B_{\mathbf{D}}, Z \geq Y),$$

whence $I(Y) = \widehat{\overline{\rho}}(Y)$, as needed to be shown.

The final assertion of Theorem 4 follows directly from the following straightforward Lemma.

Lemma 4 *For any* $Y \in B(\Sigma, [0, 1]), A \in [Y]$, *and* $\pi \in \Pi^0 : \pi(A) = E_\pi Y$.

□

Proof of Proposition 3.

We shall show that P4 implies Union and Splitting Invariance; the converse follows from similar reasoning.

Consider any $A \in \Sigma$, $\alpha, \beta \in [0,1]$ such that $\alpha + \beta \leq 1$, and $A' \in \Lambda^0_A$ as well as $B \in \Lambda^0_A \cap \Lambda^0_{A^c}$ disjoint from A' such that $\pi^0(A'/A) = \alpha$ and $\pi^0(B/A) = \pi^0(B/A^c) = \beta$. To establish Union and Splitting Invariance, it clearly suffices to show that $\rho(A' + B) = \alpha\rho(A) + \beta$.

Pick consequences x_1, x_2 such that $u(x_1) = \beta$ and $u(x_2) = \alpha + \beta$. By utility sophistication and the induced conditional linearity property of $\widehat{\rho}$ stated in the proof of Theorem 4, $[\overline{\overline{x}}, A' + B; \overline{x}, (A' + B)^c] \sim [x_2, A; x_1, A^c]$. Moreover, exploiting P4, $[x_2, A; x_1, A^c] \sim [x_2, T; x_1, T^c]$ for any $T \in \Lambda^0$ with $\rho(T) = \rho(A)$. Hence by transitivity $\rho(A' + B) = I(1_{A'+B}) = I(u \circ [x_2, T; x_1, T^c]) = E_\rho(u \circ [x_2, T; x_1, T^c]) =$

$(\alpha + \beta)\rho(T) + \beta\rho(T^c) = \alpha\rho(A) + \beta.$ □

Proof of Proposition 4.

By Theorem 3, part 3, it is clear that $\Pi \supseteq \Pi^*$ for any $\Pi \in \mathcal{K}(\Delta(\Sigma))$ such that

$$f \succsim g \text{ whenever } E_\pi(u \circ f) \geq E_\pi(u \circ g) \text{ for all } \pi \in \Pi. \tag{16}$$

Thus one needs to show that indeed

$$f \succsim g \text{ whenever } E_\pi(u \circ f) \geq E_\pi(u \circ g) \text{ for all } \pi \in \Pi^*.$$

To verify this, consider any $A \in [u \circ f]$ and $B \in [u \circ g]$. By utility sophistication, $f \succsim g$ if and only if $A \geq B$. In view of Lemma 3, if $E_\pi(u \circ f) \geq E_\pi(u \circ g)$ for all $\pi \in \Pi^*$, then also $\pi(A) \geq \pi(B)$ for all $\pi \in \Pi^*$, and therefore $A \geq B$. Since by utility sophistication, $f \succsim g$ if and only if $A \geq B$, one obtains $f \succsim g$ as desired. □

Proof of Proposition 8.

By Proposition 4, $\Pi_{\left(\geq_{\mathcal{M}(\Sigma_{obs})}\right)^*}$ is the smallest set $\Pi \in \mathcal{K}(\Delta(\mathcal{M}(\Sigma_{obs})))$ such that, for all $Y, Z \in B(\mathcal{M}(\Sigma_{obs}), [0,1])$,

$$\widehat{\rho}(Y) \geq \widehat{\rho}(Z) \text{ whenever } E_\pi Y \geq E_\pi Z \text{ for all } \pi \in \Pi. \tag{17}$$

Since I equals the restriction of $\widehat{\rho}$ to Σ_{obs}, clearly for all $Y, Z \in B(\mathcal{M}(\Sigma_{obs}), [0,1])$,

$$I(Y) \geq I(Z) \text{ whenever } E_\pi Y \geq E_\pi Z \text{ for all } \pi \in \Pi^{us}_{\succsim_{obs}}.$$

Hence

$$f \succsim_{obs} g \text{ whenever } E_\pi f \geq E_\pi g \text{ for all } \pi \in \Pi^{us}_{\succsim_{obs}}.$$

Consider any set $\Pi' \in \mathcal{K}(\Delta(\Sigma_{obs}))$ such that

$$I(Y) \geq I(Z) \text{ whenever } E_\pi Y \geq E_\pi Z \text{ for all } \pi \in \Pi'.$$

For any $\pi \in \Delta(\Sigma_{obs})$, define its extension to $\mathcal{M}(\Sigma_{obs})$ $\widetilde{\pi}$ by setting

$$\widetilde{\pi}(A) = \sum_{S \in \mathcal{S}} \pi(S)\, \pi^0(A \cap S / S),$$

and define $\widetilde{\Pi} : \mathcal{K}(\Delta(\mathcal{M}(\Sigma_{obs})))$ as $\{\widetilde{\pi} \mid \pi \in \Pi'\}$.

We shall verify that $\widetilde{\Pi}$ satisfies (17), from which it follows that $\widetilde{\Pi} \supseteq \Pi_{\left(\geq_{\mathcal{M}(\Sigma_{obs})}\right)^*}$, hence that $\Pi' = rest_{\Sigma_{obs}} \widetilde{\Pi} \supseteq rest_{\Sigma_{obs}} \Pi_{\left(\geq_{\mathcal{M}(\Sigma_{obs})}\right)^*} = \Pi^{us}_{\succsim_{obs}}$, as needs to be shown.

It remains to verify that, for all $Y, Z \in B(\mathcal{M}(\Sigma_{obs}), [0,1])$,

$$\widehat{\rho}(Y) \geq \widehat{\rho}(Z) \text{ whenever } E_{\widetilde{\pi}} Y \geq E_{\widetilde{\pi}} Z \text{ for all } \widetilde{\pi} \in \widetilde{\Pi}.$$

But this is a straightforward from noting that, by Conditional Linearity of $\widehat{\rho}$ (cf. Lemma 2), $\widehat{\rho}(Y) = \widehat{\rho}\left(E^0(Y \mid \mathcal{S})\right) = I\left(E^0(Y \mid \mathcal{S})\right)$, as well as the identity $E_{\widetilde{\pi}} Y = E_{\widetilde{\pi}} E^0(Y \mid \mathcal{S}) = E_\pi E^0(Y \mid \mathcal{S})$. □

Proof of Proposition 9. Omitted, since it parallels the first part of the proof of Theorem 3.

Proof of Proposition 10.

The equivalence of (1) and (2) is immediate from the definitions. In view of Proposition 9, (2) is clearly equivalent to the following condition on I. For all $\alpha \in (0,1]$ and all $f \in B$,

$$I(\alpha 1_A + (1-\alpha) f) = I(\alpha I(1_A)\mathbf{1} + (1-\alpha) f) = I(\alpha 1_A) + I((1-\alpha) f);$$

the second equality follows from $c-$linearity.

By straightforward but slightly tedious manipulation, this condition in turn is equivalent to the following. For all f, g and f', g' such that $f - g$ is $\{A, A^c\}$-measurable and $f - g = f' - g'$,

$$I(f) - I(g) = I(f') - I(g'). \tag{18}$$

Equation (18) clearly implies part (3) of the Proposition. Conversely, suppose that (18) is violated, i.e. that there exist f, g and f', g' such that $f - g$ is $\{A, A^c\}$-measurable and $f - g = f' - g'$, and $\varepsilon > 0$ such that

$$I(f) - I(g) > \varepsilon > I(f') - I(g').$$

By construction, the range of each of the four RVs $\frac{1}{2}f + \frac{1}{4}, \frac{1}{2}f' + \frac{1}{4}, \frac{1}{2}(g + \varepsilon) + \frac{1}{4}, \frac{1}{2}(g' + \varepsilon) + \frac{1}{4}$ is contained in $[0, 1]$. Hence by c-linearity,

$$I(\frac{1}{2}f + \frac{1}{4}) - I(\frac{1}{2}(g + \varepsilon\mathbf{1}) > 0 > I(\frac{1}{2}f' + \frac{1}{4}) - I(\frac{1}{2}(g' + \varepsilon\mathbf{1}). \tag{19}$$

But this means that $\frac{1}{2}f + \frac{1}{4} \succ_{obs} \frac{1}{2}(g + \varepsilon\mathbf{1})$ as well as $\frac{1}{2}f' + \frac{1}{4} \prec_{obs} \frac{1}{2}(g' + \varepsilon\mathbf{1})$, in contradiction to part (3). $\square$

REFERENCES

[1] Allais, M. (1953): "Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Américaine", *Econometrica* 21, 503-546.

[2] Anscombe, F. J. and R. J. Aumann (1963): "A Definition of Subjective Probability," *Annals of Mathematical Statistics*, 34, pp. 199-205.

[3] Bewley, T. F. (1986): "Knightian Decision Theory, Part I," Cowles Foundation Discussion Paper No. 807.

[4] Casadesus, R. , P. Klibanoff and E. Ozdenoren: "Maxmin Expected Utility over Savage Acts with a Set of Priors", *Journal of Economic Theory* 92, 35-65.

[5] Choquet, G. (1953): "Theory of Capacities", *Ann. Instit. Fourier* (Grenoble) 5, 131-295.

[6] Dempster, A. (1967), "Upper and Lower Probabilities Induced by a Multi-Valued Mapping", *Annals of Mathematical Statistics* 38, 325-339.

[7] Eichberger, J. and D. Kelsey (1996): "Uncertainty-Aversion and Preference for Randomisation", *Journal of Economic Theory 71,* 31-43.

[8] Ellsberg, D. (1961): "Risk, Ambiguity, and the Savage Axioms", *Quarterly Journal of Economics* 75, 643-669.

[9] Epstein, L. (1999): "A Definition of Uncertainty Aversion", *Review of Economic Studies* 66, 579-608.

[10] Epstein, L. and J.-K. Zhang (2001): "Subjective Probabilities on Subjectively Unambiguous Events", *Econometrica* 69, 265-306.

[11] Gilboa, I. and D. Schmeidler (1989): "Maxmin Expected Utility with a Non-Unique Prior", *Journal of Mathematical Economics* 18, 141-153.

[12] Ghirardato, P. and M. Marinacci (2001a): "Ambiguity Made Precise: A Comparative Foundation", *Journal of Economic Theory,* forthcoming.

[13] Ghirardato, P. and M. Marinacci (2001b): "Risk, Ambiguity, and the Separation of Utility and Beliefs", *Mathematics of Operations Research,* forthcoming.

[14] Ghirardato, P., F. Maccheroni and M. Marinacci (2001c): "A Fully Subjective Perspective on Ambiguity," in progress October 2001.

[15] Ghirardato, P., F. Maccheroni, M. Marinacci and M. Siniscalchi (2001d): "Subjective Foundations for Objective Randomization: A New Spin on Roulette Wheels", mimeo.

[16] Kannai, Y. (1992): "The Core and Balancedness", in: R. Aumann and S. Hart (eds.), *Handbook of Game Theory,* North Holland, Amsterdam, 355-395.

[17] Klibanoff, P. (2001a): "Stochastically Independent Randomization and Uncertainty Aversion", *Economic Theory* 18, 605-620.

[18] Klibanoff, P. (2001b): "Characterizing Uncertainty Aversion Through Preference for Mixtures", *Social Choice and Welfare* 18, 289-301.

[19] Kraft, C. , Pratt, J. and A. Seidenberg (1959): "Intuitive Probability on Finite Sets," *Annals of Mathematical Statistics* 30, 408-419.

[20] Machina, M. and D. Schmeidler (1992): "A More Robust Definition of Subjective Probability", *Econometrica* 60, 745-780.

[21] Nehring, K. (1991): *A Theory of Rational Decision with Vague Beliefs.* Ph.D. dissertation, Harvard University.

[22] Nehring, K. (1994): "On the Interpretation of Sarin and Wakker's 'A Simple Axiomatization of Nonadditive Expected Utility Theory"', *Econometrica* 62, 935-938.

[23] Nehring, K. (1996): "Preference and Belief without the Independence Axiom", talk presented at LOFT2 in Torino, Italy.

[24] Nehring, K. (1999): "Capacities and Probabilistic Beliefs: A Precarious Coexistence", *Mathematical Social Sciences* 38, 197-213.

[25] Nehring, K. (2000): "Rational Choice under Ignorance", *Theory and Decision* 48, 205-240.

[26] Nehring, K. (2001): "Imprecise Qualitative Probability", in progress.

[27] Ramsey, F. (1931): "Truth and Probability", in *The Foundations of Mathematics and other Logical Essays,* reprinted in: H. Kyburg and H. Smokler (eds., 1964), *Studies in Subjective Probability,* Wiley, New York, 61-92.

[28] Sarin, R. and P. Wakker (1992): "A Simple Axiomatization of Nonadditive Expected Utility Theory", *Econometrica* 60, 1255-1272.

[29] Sarin, R. and P. Wakker (1995): "On the Interpretation of Likelihood in Choquet Expected Utility", mimeo, UCLA and University of Leiden.

[30] Savage, L.J. (1954). *The Foundations of Statistics.* New York: Wiley. Second edition 1972, Dover.

[31] Schmeidler, D. (1989): "Subjective Probability and Expected Utility without Additivity", *Econometrica 57,* 571-587.

[32] Smith, C. (1961): "Consistency in Statistical Inference and Decision," *Journal of the Royal Statistical Society*, Series B, 22, pp. 1-25.

[33] Shafer, G. (1976). *A Mathematical Theory of Evidence*, Princeton: Princeton U.P.

[34] Wakker, P. (1989): *Additive Representations of Preferences.* Dordrecht: Kluwer.

[35] Wakker, P. (2001): "On the Composition of Risk Preferences and Belief", *Psychological Review*, forthcoming.

[36] Walley, P. (1991): *Statistical Reasoning with Imprecise Probabilities*. London: Chapman and Hall.

[37] Zhang, J.-K. (1997): "Subjective Ambiguity, Probability and Capacity", mimeo, University of Toronto.