Ellsberg Without Allais: A Theory of Utility-Sophisticated Preferences under Ambiguity

Klaus Nehring
University of California, Davis, and
Institute for Advanced Study, Princeton*

December 2003†

ABSTRACT

A decision-maker is utility-sophisticated if he ranks acts according to their expected utility, whenever such comparisons are well-defined in light of his potential probabilistic beliefs. We characterize utility sophistication in cases in which probabilistic beliefs are not too imprecise (i.e. "convex-ranged"), and show that in those cases preferences are completely determined by consequence utilities and event attitudes captured by preferences over bets. The Anscombe-Aumann framework as employed in the classical contributions by Schmeidler (1989) and Gilboa-Schmeidler (1989) can be viewed as an important special case.

*e-mail: kdnehring@ucdavis.edu

†Some of the material of this paper is contained in an earlier working paper "Ambiguity in the Context of Probabilistic Beliefs" and were presented at RUD 2002 in Paris.

1. INTRODUCTION

Expected Utility theory rests on two pillars of consequentialist rationality: the existence of a unique subjective probability measure underlying all decisions, and the linear determination of preferences over multi-valued acts by cardinal utilities. Descriptively, both of these assumptions have been challenged: on the one hand, it is frequently not possible to represent the betting preferences of empirical decision-makers in terms of a well-defined subjective probability measure as in the Ellsberg paradox; in such cases, decision-makers are said to view certain events as "ambiguous". On the other hand, faced with given probabilities, utilities and probabilities may not combine linearly, as in the Allais paradox and related phenoma; such decision-makers are sometimes referred to as exhibiting "probabilistic risk-attitudes".

While a descriptively fully adequate model of decision-making (if such a model exists!) will need to incorporate both phenomena, for modelling purposes it will often be desirable to zoom in on one of these two central departures from the expected utility paradigm. To this purpose, Machina-Schmeidler (1992) have introduced the notion of probabilistic sophistication, which precludes all phenomena of ambiguity but does not constrain the nature of probabilistic risk-attitudes. In the present paper, we shall introduce a "dual" notion of utility sophistication which precludes all phenomena of probabilistic risk-attitudes but does not constrain the decision-maker's attitudes towards ambiguity.

This notion of utility sophistication has also a distinct normative motivation. Since the two departures from expected utility are fundamentally different in kind, it can be argued that one departure is rationally justifiable while the other is not. Utility Sophistication in particular captures the normative position that departures due to ambiguity are justifiable while those due to probabilistic risk-aversion are not (at a fundamental level). Indeed, it can even be argued that rational decision making under partial or complete ignorance cannot rationally be based well-defined subjective probabilities (see the classical literature on complete ignorance surveyed in Luce-Raiffa (1957) as well the subsequent contributions of Jaffray (1989) and Nehring (1991,2000)). By contrast, we are not aware of an argument that would rationally force departures from expected utility in the presence of probabilities (under risk); moreover, the typical examples of such departures can be naturally accounted for either in terms of the existence of implicit psychological payoffs (cf. Broome (1991)) or as cognitive distortions in the processing of probabilities (Wakker). This paper does not attempt to defend this normative view but simply tries to articulate it axiomatically.

Broadly speaking, we shall view an agent as "utility sophisticated" if he compares acts in terms of their expected utility "whenever possible". Since the possibility of such a comparison depends on the agents' beliefs, utility sophistication must be defined relative to a specified set of beliefs; we shall model beliefs as partial orderings over events (likelihood relations) represented by a (closed convex) set of admissible probability measures Π. The agent's beliefs might, in principle, be derived from his preferences; however, it seems doubtful that this can be done in a canonical way in general, analogously to Savage's canonical definition of revealed likelihood (see Nehring 2001 for a detailed analysis). In order to bracket this issue for the bulk of the analysis, we shall therefore assume that (some of) the decision-maker's beliefs are independently given; at the end of the paper, we investigate to what extent this assumption can be weakened or given up. Given a cardinal utility function u (obtained from preferences), an agent is *utility sophisticated* with respect to the set of admissible priors Π if the agent prefers any act f over another act g whenever the expected utility of f weakly exceeds that of g with respect to any admissible prior.

Note that utility sophistication implies expected utility maximization over unambiguous acts (acts whose induced distribution does not depend on the prior), but is substantially stronger than the latter by restricting preferences over certain comparisons among ambiguous acts. This added strength proves crucial for the analytical power of utility sophistication.

The first and foremost task of the paper is to provide axiomatic foundations. The crucial axiom replacing Savage's Sure-Thing Principle is the axiom of "Unambiguous Tradeoff Consistency"; its basic idea can be described as follows: suppose that we already have obtained a ranking of utility-differences from the decision-maker's preferences over unambiguous acts, and consider two acts f and g whose outcomes differ on only two *equally likely* events A and B such with f yielding a better outcome in A and g in B. Unambiguous Tradeoff Consistency requires that if the utility gain from f over g in A exceeds the utility gain from g over f in B, then f is preferred to g. In order to derive utility sophistication from this axiom, we need to assume both a rich set of consequences and a sufficiently precise (specifically: "convex-ranged") likelihood relation. As an important additional benefit, these assumptions imply that preferences are *bi-determinate*, that is: determined by preferences over two-outcome acts only; indeed, somewhat more specifically, preferences over general acts are determined by preferences over bets (which capture the decision-maker's perceived ambiguity and ambiguity attitudes) on the one hand, and preferences over two-outcome unambiguous acts (which reflect the decision-maker's cardinal valuation of consequences). In this, utility-sophisticated preferences are analogous to Choquet expected utility preferences which are also bi-determinate. However, utility-sophistication and Choquet expected utility are inherently different ways of constructing multi-outcome preferences from two-outcome ones; they agree only in the degenerate case of subjective expected utility (Proposition 4).

Bi-determinacy implies that any further assumptions on preferences can be broken down in assumptions on risk-attitudes on the one hand and on ambiguity-attitudes on the other, the latter via assumptions on betting preferences. One basic condition of this kind is the requirement that betting preferences do not depend on stakes (winning and losing outcomes) involved, i.e. Savage's axiom P4. In Theorem 2, we show that under utility-sophistication, P4 is equivalent to the "constant-linearity" of the evaluation functional, a condition that plays a fundamental role in the work of Ghirardato and Marinacci. Further natural conditions describe the ambiguity-valorization (aversion versus seeking) in terms of betting preferences; these are explored in the companion paper Nehring (2004) and yield a novel and arguably improved axiomatization of the Minimum Expected Utility model due to Gilboa-Schmeidler (1989).

A basic instance of convex-ranged beliefs is derived from the existence of an independent continuous random device by modelling the random device as part of the state space rather than the outcome space, as in the traditional Anscombe-Aumann (1963) framework. We show in Proposition 3 that the "background" assumptions contained in the classical contributions of Schmeidler (1989) and Gilboa-Schmeidler (1989) and others in the Anscombe-Aumann framework are equivalent to utility-sophistication (relative to the beliefs capturing the random device) combined with stake-independence in our approach. Thus these assumptions prove to have much more substance than meets the eye!

In this setting, a decision-maker will typically have further beliefs (represented by a more precise, i.e. smaller, set of priors). Suppose that the decision-maker is utility-sophisticated relative to his randomization beliefs, but not relative to his total set of beliefs. Since this decision-maker would sometimes rank acts contrary to the expected-utility comparisons entailed by his beliefs, assuming that he does not do so relative to his randomization beliefs, this would introduce asymmetry between two kinds of beliefs, and may engender the suspicion that the problematic distinction between objective and subjective probabilities assumed in the Anscombe-Aumann framework has sneaked in through the back door. Somewhat remarkably, however, this possibility can never arise: as shown by Proposition 5, if the decision-maker is utility-sophisticated relative to his randomization beliefs, he will be utility-sophisticated relative to his total set of beliefs, whatever they are (assuming compatibility with preferences in a natural way). This means also that, as long as one can assume that the decision-makers beliefs reflect the nature of the randomization (an assumption that can be tested by looking at his betting preferences), his utility sophistication can be ascertained from knowledge of his preferences, and knowledge of his beliefs is not required.

Suppose, however, that an outside analyst has access only to the decision-maker's preferences, and does not even know that certain events correspond to the realization of a random device with intersubjectively known probabilities. Is it still possible to determine whether the decision-maker is "truly" utility-sophisticated? Note that the analyst can still determine from preference alone that the decision-maker is consistent with the randomization beliefs, *if these were in fact his beliefs.* Now, however, the analyst has no longer an independent reason to support this assumption; it seems conceivable that she may come up with other hypothetical sets of beliefs which seem as or more "plausible" in the light of preference information alone, and relative to which the decision-maker is not utility-sophisticated. Clearly, a convincing answer can only be expected if some constraints are accepted on what beliefs can plausibly attributed to a decision-maker in light of his preference relation. To this purpose, we propose some plausibility criteria in section 6, and argue that a good case can be made to judge any decision-maker who is utility-sophisticated with respect to *some* (unobserved) sufficiently rich, i.e. convex-ranged, set of beliefs as utility-sophisticated *tout court.* We also show that this case is especially strong in a setting with a continuous randomization device and a finite set of non-random states.

While the existing literature has not yet attempted to define a notion of utility-sophistication formally, as already mentioned, a fair number of models of decision making under ambiguity in the Anscombe-Aumann framework build in utility-sophistication, starting from the seminal contributions of Schmeidler (1989) and Gilboa-Schmeidler (1989). Other contributions, especially Ghirardato-Marinacci (2002) and Ghirardato et al. (2002), assume a utility-sophisticated viewpoint by assuming in the interpretation of their definitions and axioms that all departures from expected utility can be attributed to ambiguity; Ghirardato et al. (2002) following and extending the talk Nehring (1996) define revealed beliefs assuming utility-sophistication with respect to those beliefs *by definition.*

The remainder of the paper is organized as follows. In section 2, we introduce likelihood relations and their multi-prior representation, as well as basic assumptions on preferences maintained throughout. We then define the notion of utility sophistication and characterize it axiomatically (section 3), and pay particular attention to the case of "stake-independent" (P4) betting preferences (section 4). In section 5, we study utility-sophistication in various preference models in the literature, and establish in particular a close link to standard models in the Anscombe-Aumann (1963) framework. Section 6 discusses under what conditions one can ascertain an agent's utility-sophistication with only incomplete knowledge of his beliefs, or with none at all. The concluding section 7 suggests possible generalizations and extensions. All proofs are contained in the appendix.

2. BACKGROUND

2.1. Convex-ranged Likelihood Relations

Since utility sophistication is defined relative to a specified set of probabilistic beliefs, we shall model a decision maker in terms of two entities, a preference relation $\succsim$ over Savage acts and a comparative likelihood relation $\trianglerighteq$ describing his probabilistic beliefs. Formally, a likelihood relation is a partial ordering $\trianglerighteq$ on an algebra of events Σ in a state space Ω, 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 $\equiv$. We shall treat the likelihood relation as an independent primitive, leaving open the possibility that it may be derived from the decision maker's preferences themselves. Alternatively, the likelihood relation may be viewed as a "parameter" in expressions such as "is utility sophisticated relative to $\trianglerighteq$", which renders such as expressions predicates of preferences only; on this latter interpretation, there is no presupposition that there is an independent matter of fact as to whether the decision maker truly has those beliefs $\trianglerighteq$; see Nehring (2003) for a more detailed discussion of this point and the following concepts, where this framework of "decision-making in the context of probabilistic beliefs" has been introduced.

A *prior* π is a finitely-additive set-function on Σ such that $\pi(\Omega) = 1$. Given a likelihood relation $\trianglerighteq$, let $\Pi_{\trianglerighteq}$ denote its set of *admissible priors* defined by

$$\pi \in \Pi_{\trianglerighteq} \text{ if and only if, for all } A, B \in \Sigma,\ A \trianglerighteq B \text{ whenever } \pi(A) \geq \pi(B).$$

For any $\trianglerighteq$, $\Pi_{\trianglerighteq}$ is a closed convex set in the product (or weak*) topology. Conversely, any set of priors Π induces an associated likelihood relation $\trianglerighteq_{\Pi}$ given by the unamimity condition

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

A likelihood relation $\trianglerighteq$ is *non-contradictory* if $\Pi_{\trianglerighteq} \neq \varnothing$. It is *coherent* if there exists $\Pi \neq \varnothing$ such that $\trianglerighteq = \trianglerighteq_{\Pi}$; clearly, if $\trianglerighteq = \trianglerighteq_{\Pi}$, then also $\trianglerighteq = \trianglerighteq_{\Pi'}$ where Π' is the closed convex hull of Π; it is therefore without loss of generality to assume sets of priors to be closed and convex. It is also easily verified that $\trianglerighteq$ is coherent if and and only if $\trianglerighteq = \trianglerighteq_{(\Pi_{\trianglerighteq})}$; the set $\Pi_{\trianglerighteq}$ will therefore be referred to as the multi-prior representation of $\trianglerighteq$. Note, however, that in general, the multi-prior representation may not be unique in that there may exist closed convex sets $\Pi' \neq \Pi_{\trianglerighteq}$ such that $\trianglerighteq = \trianglerighteq_{\Pi'}$; in those cases in which , there is a "loss of information" in representing beliefs by likelihood relations rather than sets of priors.

A central role in the following will be played by likelihood relations with a convex-ranged multi-prior representation. The set of priors Π is *convex-ranged* if, for any event $A \in \Sigma$ and any $\alpha \in (0,1)$, there exists an event $B \in \Sigma$, $B \subseteq A$ such that $\pi(B) = \alpha\pi(A)$ for all $\pi \in \Pi$; we shall refer to the associated likelihood relation $\trianglerighteq_{\Pi}$ likewise as convex-ranged. Nehring (2003) contains an axiomatization of coherent convex-ranged likelihood relations, and shows that their multi-prior representation is always unique; the latter implies that, for any convex-ranged Π, $\Pi = \Pi_{(\trianglerighteq_{\Pi})}$. Convex-ranged likelihood relations are characterized by a rich set of unambiguous and conditionally unambiguous events. Say that B is *unambiguous given* A if, for some $\alpha \in [0,1]$, $\pi(B) = \alpha\pi(A)$ for all $\pi \in \Pi_{\trianglerighteq}$. Let Λ_A denote the family of events that are unambiguous given A; clearly, A is closed under finite disjoint union and under complementation in A, but not necessarily under intersection. An event A is $\trianglerighteq$-*null* if $A \equiv \varnothing$, or, equivalently, if $\pi(A) = 0$ for all $\pi \in \Pi_{\trianglerighteq}$. For any non-null A, let $\overline{\pi}(./A)$ denote the restriction of $\pi(./A)$ to Λ_A for any $\pi \in \Pi_{\trianglerighteq}$. We will say that B is *unambigous* if it is "unambiguous given Ω", and write Λ for Λ_{Ω}, as well as $\overline{\pi}$ for $\overline{\pi}(./\Omega)$. It is easily verified that a coherent likelihood relation $\trianglerighteq$ defined on a $\sigma-$algebra is convex-ranged if and only if, for all events $A \in \Sigma$, there exists an event $B \in \Sigma$ such that $B \subseteq A$ and $B \equiv A \backslash B$. Thus convex-rangedness assumes that the likelihood relation is not "too incomplete". As a matter of further notion, let $\pi^{-}(A) = \min_{\pi \in \Pi} \pi(A)$ and $\pi^{+}(A) = \max_{\pi \in \Pi} \pi(A)$ the lower and upper probabilities of event A.

Example 1 (Continuous Randomization Device). Consider a state space that can be written as $\Omega = \Omega_1 \times \Omega_2$, where the space Ω_1 is the space of "generic states" , and Ω_2 that of independent "random states" with associated algebras Σ_1 and algebra Σ_2. Let η denote a convex-ranged[1], finitely additive prior over random events Σ_2 . The "continuity" and stochastic independence of the random device are captured by the following coherent likelihood relation $\trianglerighteq_{AA}$ defined on the product algebra $\Sigma = \Sigma_1 \times \Sigma_2$ noting that any $A \in \Sigma_1 \times \Sigma_2$ can be written as $A = \sum_i S_i \times T_i$, where the $\{S_i\}$ form

[1]that is, $\{\eta\}$ is convex-ranged in Σ_2.

a finite partition of Ω_1 :[2]

$$\sum_i S_i \times T_i \trianglerighteq_{AA} \sum_i S_i \times T_i' \text{ if and only if } \eta(T_i) \geq \eta(T_i') \text{ for all } i.$$

Clearly, there exists a unique set of priors Π_{AA} representing $\trianglerighteq_{AA}$; indeed, Π_{AA} is simply the set of all product-measures $\pi_1 \times \eta$ (where π_1 ranges over all finitely additive measures on Σ_1), reflecting the stochastic independence of the random device. Note that the convex-rangedness of Π_{AA} is a straightforward consequence of the convex-rangedness of η.

In general, a decision-maker will have further beliefs in this setting captured by a likelihood relation $\trianglerighteq$ containing $\trianglerighteq_{AA}$; this relation evidently inherits the convex-rangedness of $\trianglerighteq_{AA}$.

Example 2 (Limited Imprecision). A particular way to make the intuition of a limited extent of overall ambiguity precise is to assume that Σ is a σ-algebra and that Π is the convex hull of a *finite* set Π' of *non-atomic, countably additive* priors. Due to Lyapunov's (1940) celebrated convexity theorem, Π is convex-ranged. The priors $\pi \in \Pi'$ can be interpreted as a finite set of hypotheses a decision-maker deems reasonable without being willing to assign probabilities to them. Finitely generated sets of priors also occur naturally in social belief aggregation, where $\trianglerighteq_I$ represents the unanimity likelihood ordering induced by the finite set of individuals' likelihood orderings $\trianglerighteq_i$ that are assumed to be precise with representing measures μ_i.

In the following, we shall on occasion make use of the strict component of a coherent likelihood relation; specifically, let $A \vartriangleright B$ ("A is strictly more likely than B") if $\pi(A) > \pi(B)$ for all $\pi \in \Pi_{\trianglerighteq}$. Note that, in general, $\vartriangleright$ is a proper subrelation of the asymmetric component of $\trianglerighteq$; also, by the compactness of $\Pi_{\trianglerighteq}$, $A \vartriangleright B$ if and only if $\min_{\pi\in\Pi} \pi(A) - \pi(B) > 0$. Under convex-rangedness, $\vartriangleright$ can be characterized in terms of $\trianglerighteq$ as follows.

Fact 1 *Suppose that $\trianglerighteq$ is a convex-ranged belief context. Then $A \vartriangleright B$ if and only if there exist events $A' \subseteq A$, $B' \subseteq B$ and $T \in \Lambda$ disjoint from A' such that $A' \equiv A \backslash A'$, $B' \equiv B \backslash B'$, not $T \equiv \emptyset$ and $A' \trianglerighteq B' + T$.*

2.2 Structural Assumptions on Preferences

Consider now a DM described by a preference ordering over acts and a coherent likelihood relation $\trianglerighteq$ that will also referred to as the DM's (belief) context; we will typically write Π for $\Pi_{\trianglerighteq}$. Let X

[2]The relation $\trianglerighteq_{AA}$ is easily characterized axiomatically; see Nehring (2003).

be a set of *consequences*. An *act* is a finite-valued mapping from states to consequences, $f : \Omega \rightarrow X$, that is measurable with respect to the 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 $[x_1$ on $A_1; x_2$ on $A_2; ...]$ for the act with consequence x_i in event A_i; constant acts $[x$ on $\Omega]$ are typically referred to by their constant consequence x. To prepare the ground for the subsequent analysis, we now introduce the basic substantive and regularity assumptions that will be maintained throughout.

Explicit beliefs determine most directly preferences over bets. A *bet* is a pair of acts with the same two outcomes, i.e. a pair of the form $([x \text{ on } A; y \text{ on } A^c], [x \text{ on } B; y \text{ on } B^c])$. Fundamental is the following rationality requirement on the relation between preferences and probabilistic beliefs.[3]

Axiom 1 (Compatibility) *For all $A, B \in \Sigma$ and $x, y \in X$:*

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

Preferences will be assumed to be "state-independent" throughout; in particular, they will be assumed to be eventwise monotone in the following weak version of Savage's axiom P3.

Axiom 2 (Eventwise Monotonicity) *For all acts $f \in \mathcal{F}$, consequences $x, y \in X$ and events $A \in \Sigma$: $[x$ on $A; f(\omega)$ elsewhere$] \succsim [y$ on $A; f(\omega)$ elsewhere$]$ whenever $x \succsim y$.*

The following condition ensures in particular that the set of consequences is sufficiently rich.

Axiom 3 (Solvability) *For any $x, y \in X$ and $T \in \Lambda$, there exists $z \in X$ such that $z \sim [x, T; y, T^c]$.*

For expositional simplicity especially in the stake-dependent case, we shall assume throughout that the preference relation is bounded in utility.

Axiom 4 (Boundedness) *There exist x^- and $x^+ \in X$ such that, for all $x \in X$, $x^- \precsim x \precsim x^+$.*

To obtain a real-valued representation, some Archimedean property is usually assumed. The following is sufficiently strong to help deliver the main result, Theorem 1, below. Note that it is defined relative to explicit beliefs, and presumes their convex-rangedness. Substantively, as confirmed

[3]The use in the second clause of the strict likelihood relation $\triangleright$ rather than of the asymmetric component of $\trianglerighteq$ is necessary in order to preserve generality; for example, if bets are ranked according to the lower probabilty of the superior outcome x $\pi^-(\{\omega : f(\omega) = x\})$, there may exist events A with $\pi^+(A) > 0$ and $\pi^-(A) = 0$; in this case $[x$ on $A; y$ on $A^c] \sim [x$ on $\emptyset; y$ on $\Omega]$ even though $A \triangleright^0 \emptyset$.

by the upcoming representation result, Proposition 1, it asserts that if acts are changed on events of sufficiently small (upper) probability, strict preference does not change.

Axiom 5 (Archimedean) *For any acts $f = [x$ on A, y on B; f otherwise$]$ and g such that $f \succ g$ (resp. $f \prec g$) and such that A is unambiguous given $A+B$, there exists an event that is unambiguous given $A+B$ such that $A' \lhd A$ and $f' = [x$ on A', y on $A+B-A'$; f otherwise$] \succ g$ (resp. $f' \prec g$).*

Let $\mathcal{Z}$ denote the set of finite-valued, Σ-measurable functions $Z : \Omega \to [0,1]$. Under the above axioms, we will now establish a basic representation theorem that ensures the existence of a utility function u mapping X onto the unit interval and of an evaulation functional I: $\mathcal{Z} \to [0,1]$ such that $f \succsim g$ if and only if $I(u \circ f) \geq I(u \circ g)$, for all $f, g \in \mathcal{F}$.

I is *normalized* if $I(c\mathbf{1}) = c$ for all $c \in [0,1]$ and $I(1_T) = \overline{\pi}(T)$ for all $T \in \Lambda$. Note that for normalized I, by compatibility, $I(1_A) = \rho(A)$ for all $A \in \Sigma$; moreover, u must be calibrated in terms of probabilities as well, i.e. satisfy $u(z) = \overline{\pi}(T)$ whenever $z \sim [x^+, T; x^-, T^c]$.[4] I is *monotone* if $I(Y) \geq I(Z)$ whenever $Y \geq Z$; I is *compatible with* $\trianglerighteq$ if $I(1_A) \geq I(1_B)$ whenever $A \trianglerighteq B$ and $I(1_A) > I(1_B)$ whenever $A \rhd B$; I is *event-continuous* if, for any $x, y \in X$, $Z \in \mathcal{Z}$, $E \in \Sigma$, $A \in \Lambda_E$ and any increasing sequence $\{A_n\}$ of events contained in A such that $\overline{\pi}(A_n/E)$ converges to $\overline{\pi}(A/E)$, $I(x1_{A_n} + y1_{E \setminus A_n} + Z1_{E^c})$ converges to $I(x1_A + y1_{E \setminus A} + Z1_{E^c})$.

Proposition 1 *Let $\trianglerighteq$ be a convex-ranged belief context. The following two statements are equivalent:*

i) the preference ordering $\succsim$ is compatible with $\trianglerighteq$, Archimedean, solvable, bounded and eventwise monotone.

ii) there exist an onto utility-function $u : [0,1]$ and a functional $I : \mathcal{Z} \to [0,1]$ that is monotone, event-continuous and compatible with $\trianglerighteq$ such that

$$f \succsim g \text{ if and only if } I(u \circ f) \geq I(u \circ g), \text{ for all } f, g \in \mathcal{F}.$$

There is a unique pair (u, I) satisfying ii) such that I is normalized.

In the sequel, betting preferences will play a special role. We shall frequently but not always assume that preferences over bets depend only on the events involved, not on the stakes. This is captured by Savage's axiom P4.

Axiom 6 (Stake Independence, P4) *For all $x, y, x', y' \in \mathcal{X}$ such that $x \succ y$ and $x' \succ y'$ and all $A, B \in \Sigma$:*

[4]To see this, $z \sim [x^+, T; x^-, T^c]$ implies $I(u(z)\mathbf{1}) = I(1_T)$. Thus by the two normalization conditions $u(z) = I(u(z)\mathbf{1}) = I(1_T) = \overline{\pi}(T)$.

$[x \text{ on } A; y \text{ on } A^c] \succsim [x \text{ on } B; y \text{ on } B^c]$ *iff* $[x' \text{ on } A; y' \text{ on } A^c] \succsim [x' \text{ on } B; y' \text{ on } B^c]$.

We will use the notation $A \succsim B$ and for the preference $[x^+ \text{ on } A; x^- \text{ on } A^c] \succsim [x^+ \text{ on } B; x^- \text{ on } B^c]$, and will denote the entire event-preference relation $\succsim$ on Σ by $\succsim_{bet}$ to avoid confusion with the act-preference relation. This notation is primarily motivated by the stake-independent case in which the event-preference relation $\succsim_{bet}$ completely summarizes the DM's beliefs and ambiguity attitudes.[5] If preferences are utility-sophisticated, this will be the case even when betting preferences are stake-dependent.

Compatibility of betting preferences with explicit beliefs ensures a ranking of bets on unambiguous events according their explicit probability $\overline{\pi}$. Under the assumptions of Proposition 1, there exists a unique set-function $\rho : \Sigma \to [0,1]$ representing event-preferences $\succsim_{bet}$ that is additive on unambiguous events and has $\rho(\Omega) = 1$; ρ assigns to each event the probability $\overline{\pi}(T)$ of any unambigous event to which it is indifferent. If I is normalized, clearly $\rho(A) = I(1_A)$. ρ is *compatible with* $\trianglerighteq$ if $\rho(A) \geq \rho(B)$ whenever $A \trianglerighteq B$ and $\rho(A) > \rho(B)$ whenever $A \triangleright B$; finally, ρ is *event-continuous* if, for any disjoint $B, E \in \Sigma$, $A \in \Lambda_E$ and any increasing (respectively decreasing) sequence $\{A_n\}$ of events contained in (resp. containing) A such that $\overline{\pi}(A_n/E)$ converges to $\overline{\pi}(A_n/E)$, $\rho(A_n + B)$ converges to $\rho(A + B)$.

3. UTILITY SOPHISTICATED PREFERENCES

The fundamental goal of this paper is to provide axiomatic foundations for the intuitive notion of a decision-maker who departs from expected-utility *only* for reasons of ambiguity. This idea can be formulated transparently with reference to exogenously specified belief context $\trianglerighteq$ as described in the previous section in terms of the following property of utility sophistication. We will show later that, under certain conditions, in an important sense it is possible to overcome this reliance on explicit beliefs as a non-preference datum.

Definition 1 (Utility Sophistication) *The preference relation $\succsim$ is utility-sophisticated with respect to Π if there exists $u : X \to \mathbf{R}$ such that $f \succsim g$ (resp. $f \succ g$) whenever $E_\pi u \circ f \geq E_\pi u \circ g$ (resp. $E_\pi u \circ f > E_\pi u \circ g$) for all $\pi \in \Pi$; $\succsim$ is utility-sophisticated with respect to the context $\trianglerighteq$ if it is utility-sophisticated with respect to $\Pi_{\trianglerighteq}$.*

[5] In the stake-dependent case, $\succsim_{bet}$ is formally equivalent to Savage's "revealed likelihood" relation; such an interpretation is not warranted, however, in the presence of ambiguity, since the relation incorporates not merely beliefs in this case (however construed) but also ambiguity attitudes.

By the uniqueness of the multi-prior representation of convex-ranged contexts mentioned above, for convex-ranged Π, utility-sophistication with respect to Π is the same as utility-sophistication with respect to $\trianglerighteq_\Pi$. However, without uniqueness, it may be that $\Pi \subsetneq \Pi_{(\trianglerighteq_\Pi)}$, so that utility-sophistication with respect to a set of priors cannot be equated to utility-sophistication with respect to the associated (or any other) likelihood relation.

To motivate the key axiom underlying utility sophistication, consider first the ranking of unambiguous acts for which utility sophistication boils down to EU maximization with respect to the prior $\overline{\pi}$. Specifically, consider choices among unambiguous acts f and g with two outcomes, each of which has subjective probability one half, and assume that $f = [x \text{ on } A; y \text{ on } A^c]$ and $g = [x' \text{ on } A; y'$ on $A^c]$ with $x \succ x'$, $y' \succ y$ and $A \equiv A^c$. 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 utility loss of obtaining y rather than y'. The preference of f over g by a DM committed to this principle reveals a greater utility gain from x over x' than from y' over y. Thus, if the DM chooses $f = [x \text{ on } A; y \text{ on } A^c]$ over $g = [x' \text{ on } A; y' \text{ on } A^c]$, consistency requires that he also choose the act $[x \text{ on } E; y \text{ on } E^c]$ over $[x' \text{ on } E; y' \text{ on } E^c]$, where E is any other event that is equally likely to its complement, $E \equiv E^c$.[6] The following axiom "Unambiguous Tradeoff-Consistency" simply generalizes this consistency requirement to choices of the form $f = [x \text{ on } A; y \text{ on } B; f(\omega) \text{ elsewhere}]$ versus $g = [x' \text{ on } A; y' \text{ on } B; f(\omega) \text{ elsewhere}]$ whenever the events A and B are judged equally likely $(A \equiv B)$.

Again, also in this more general case, since the relative probabilities of the events A and B are known as equal, the comparison between the acts f and g boils down to a comparison of the respective utility gains as the only remaining relevant criterion. 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 according to 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 7 (Unambiguous Tradeoff Consistency) *For all $x, y, x', y' \in X$ such that $x \succsim x'$, acts $f, g \in \mathcal{F}$ and events A disjoint from B and A' disjoint from B' such that $A \equiv B \triangleright \emptyset$ and $A' \equiv B'$:*

[6] This consistency requirement is in fact axiom 2 of Ramsey's (1931) seminal contribution. Conditions requiring consistency of trade-offs across choices have been used elsewhere in the axiomatizations of SEU and CEU theory; see in particular Wakker (1989).

if $[x \text{ on } A; y \text{ on } B; f(\omega) \text{ elsewhere}] \succsim [x' \text{ on } A; y' \text{ on } B; f(\omega) \text{ elsewhere}]$,
then $[x \text{ on } A'; y \text{ on } B'; g(\omega) \text{ elsewhere}] \succsim [x' \text{ on } A'; y' \text{ on } B'; g(\omega) \text{ elsewhere}]$.

Note the restriction to events A and B of strictly positive lower probability; it ensures that the premise "$[x$ on $A; y$ on $B; f(\omega)$ elsewhere$] \succsim [x'$ on $A; y'$ on $B; f(\omega)$ elsewhere$]$" indeed implies that the utility difference of x over x' is not smaller than that of y' over y. Note also that Tradeoff Consistency entails Eventwise Monotonicity; indeed, Eventwise Monotonicity is simply Tradeoff Consistency restricted to cases in which $x = y$, $x' = y'$ and $A + B = \Omega$.

For Unambiguous Tradeoff Consistency (henceforth simply: "Tradeoff Consistency") to allow for ambiguity, the restriction to equally likely A and B respectively A' and B' is crucial.Indeed, if one replaced this clause by the weaker one that these events would be indifferent as bets ($A \sim B$ and $A' \sim B'$), the resulting stronger axiom would force betting preferences to satisfy the additivity condition

$$A \sim B \text{ if and only if } A + C \sim B + C, \text{ for any } A, B, \text{ and } C,$$

and thereby impose SEU. The restriction to equally likely subevents makes conceptual sense, since only then can a conditional expected utility be taken that furnishes a decisive decision criterion. In this way, Tradeoff Consistency captures the notion of a decision maker committed to expected utility principles whenver they are applicable in the presence of ambiguity.

Tradeoff Consistency becomes particularly powerful if the underlying epistemic context is convex-ranged. For not only does it entail utility sophistication, utility sophistication itself becomes particularly powerful in this case, as it implies that a DM's multi-act preferences are determined by his preferences over unambiguous acts together with his preferences over bets. Mathematically, this is the consequence of the existence of a non-linear expectation operator that is built into the DM's ambiguity attitudes, and that a plays a role somewhat similar to that of the celebrated Choquet integral.

The key to deriving this built-in expectation operator is the mixture-space structure associated with convex-ranged belief contexts introduced in Nehring (2003). With with each $Z \in \mathcal{Z}$, one can associate an equivalence class $[Z]$ of events $A \in \Sigma$ as follows. Let $A \in [Z]$ if, for there exists a partition $\{E_i\}$ of Ω such that $Z = \sum z_i 1_{E_i}$, and such that, for all $i \in I$ and $\pi \in \Pi : \pi(A \cap E_i) = z_i \pi(E_i)$. Note that $[Z]$ is non-empty by the convex-rangedness of Π. Moreover, it is easily seen that for any two $A, B \in [Z] : \pi(A) = \pi(B)$ for all $\pi \in \Pi$, and thus $A \equiv B$. Hence by Compatibility also

$A \sim_{bet} B$. One therefore arrives at a well-defined ordering of random-variables $\widehat{\succsim_{bet}}$ on $\mathcal{Z}$ by setting

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

Let $\widehat{\rho}$ denote the associated unique extension of ρ to $\mathcal{Z}$ 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

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

Clearly, by Compatibility, $\widehat{\rho}$ is a monotone, normalized evalution functional on $\mathcal{Z}$. We shall call $\widehat{\rho}(Z)$ the "intrinsic integral" of Z.

We are now in a position to state the main result of the paper.

Theorem 1 *Let $\trianglerighteq$ be a convex-ranged belief context. The following three statements are equivalent:*

1. The preference ordering $\succsim$ is Archimedean, tradeoff consistent and compatible with respect to $\trianglerighteq$, as well as bounded and solvable.

2. The preference ordering $\succsim$ is Archimedean and utility sophisticated with respect to $\trianglerighteq$, for some onto function $u : X \to [0,1]$.

3. There exists an onto function $u : X \to [0,1]$ and an event-continuous set-function ρ compatible with $\trianglerighteq$ such that, for all $f, g \in \mathcal{F}$:

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

Theorem 1 achieves two things. First of all, it delivers an axiomatic foundation for utility sophisticated preferences when the underlying belief context is convex-ranged and when the set of consequences is rich; both of these assumptions are used essentially in the derivation. As a significant surplus value, it shows that utility sophistication in the convex-ranged case implies the existence of an intrinsic-integral representation. This is important since it entails a determination of multi-act preferences by event-attitudes (captured by betting preferences and represented by ρ) and consequence-attitudes (capture by preferences over unambiguous acts and represented by u). We shall therefore call preference orderings characterized by the conditions of part three of the Theorem save event-continuity of ρ **bi-determinate utility-sophisticated.**[7]

[7]Since $\widehat{\rho}$ is constructed from ρ with reference to the context $\trianglerighteq$, it may appear that the context also plays a role in determining $\succsim$. We shall see however below that this is not the case, at least in the standard case of preferences satisfying P4.

Bi-determinacy implies that any more specific model of utility sophisticated preferences is obtained from more specific assumptions about betting preferences. Three types of assumptions come to mind in particular. First, assumptions on how betting preferences depend on the stakes involved, and among these especially Savage's "stake-independence" axiom P4. Second, assumptions on how probabilistic beliefs constrain betting preferences beyond Compatibility. And third, qualitative assumptions on the nature of a decision maker's ambiguity attitude. The impact of stake-independence will be studied in the next section, and we will see under utility-sophistication, it has close connections to assumptions of the second kind. We leave the analysis of qualitative assumptions of "ambiguity aversion" respectively "ambiguity seeking" in the present framework to a companion paper (Nehring 2004).

4. THE STAKE-INDEPENDENT CASE

A fundamental assumption on betting preferences is the independence of preferred events from the stakes involved, as expressed by Savage's axiom P4. In the context of SEU and, more generally, of probabilistic sophistication, P4 is typically viewed as a rationality axiom expressing consistency of revealed likelihood judgements. Under ambiguity, this interpretation is no longer viable in general, since betting preferences may reflect ambiguity attitudes besides likelihood judgments. On the other hand, when restricted to unambiguous events, P4 still obtains as an implication of compatibility with the underlying belief context.

We submit that, having lost its original rationale, under ambiguity P4 can no longer be viewed as a rationality condition; instead, it can be viewed as "well-behavedness" condition asserting simply that betting preferences do not dependend on the stakes involved. There does not seem to be anything strange and genuinely "inconsistent" in this. In the context of an Ellsberg urn experiment, for example, a decision maker may well prefer a bet of $1 on a draw from an urn with unknown composition (getting $0 otherwise) over a bet of $1 on an event with an objective probability of 40%, and exhibit at the same time the opposite preference once the stakes are raised to $10,000 (versus $0). Indeed, this preference pattern is naturally interpreted as reflecting greater ambiguity aversion at greater possible gains. The view of P4 as a well-behavedness rather than rationality condition allows also to resolve the conflict between dynamic consistency and P4 outside SEU observed by Epstein-Le Breton (1993).

We will now show that in the presence of utility-sophistication, P4 implies fairly substantial restrictions on the evaluation of multi-valued acts and even on betting preferences at given stakes. To state the following result, a few new terms and conditions need to be introduced. An evaluation functional I (in particular $\widehat{\rho}$) is *constant-additive* if $I(Y+c\mathbf{1})=I(Y)+c$; I is *positively homogeneous* if $I(\alpha Y)=\alpha I(Y)$ for any $\alpha\in[0,1]$; I is *constant-linear* if it is constant-additive and positively homogeneous. P4 also turns out to be equivalent to the following "additive" and "multiplicative" invariance properties of betting preferences that are of independent interest.

Axiom 8 (Union Invariance) *For any $T\in\Lambda$ and any $A,B\in\Sigma$ disjoint from T: $A\succsim B$ if and only if $A+T\succsim B+T$.*

Axiom 9 (Splitting Invariance) *For any $A,B\in\Sigma$ and any partitions of A and B into equally likely subevents $\{A_1,...,A_n\}$ and $\{B_1,...,B_n\}$,with $A_i\equiv A_j$ and $B_i\equiv B_j$ for all $i,j\leq n$, $A\succsim B$ if and only if $A_1\succsim B_1$.*

Clearly, betting preferences are union invariant if $\rho(A+T)=\rho(A)+\overline{\pi}(T)$ whenever T is unambiguous; likewise, betting preferences are splitting invariant if $\rho(B)=\overline{\pi}(B/A)\,\rho(A)$ whenver B is unambiugous given A. The following holds.

Theorem 2 *Suppose $\succsim$ is bideterminate utility sophisticated relative to the convex-ranged context $\trianglerighteq$. Then the following three statements are equivalent*

1. *$\succsim$ satisfies P4.*
2. *I is constant-linear.*
3. *$\succsim$ satisfies Union and Splitting Invariance.*

In the appendix, we derive the result by demonstrating the implications 2) $\Longrightarrow$ 1), 1) $\Longrightarrow$ 3), and 3) $\Longrightarrow$ 2). The first implication is valid for any constant-linear evaluation functional I, without reference to a convex-ranged belief context. The second implication 1) $\Longrightarrow$ 3) relates two different properties of betting preferences, making essential use of utility-sophistication. Note that in the third statement, the Union and Splitting Invariance could have been equally formulated as applying to bets at arbitrary stakes, not merely to those at extreme stakes, as has been done for neatness. Finally, the implication 3) $\Longrightarrow$ 2) mirrors the invariance properties of betting preferences in corresponding properties of the intrinsic integral $\widehat{\rho}$; bideterminate utility sophistication closes the circle via the identity $I=\widehat{\rho}$.

Preferences represented by a constant-linear evalation functional have been studied in the literature, especially by Ghirardato et al. (2002). Constant Linearity can be viewed as a cardinal stake-invariance property of multi-act preferences. Theorem 2 derives this property from the weaker and arguably more primitive ordinal P4 property, assuming utility sophistication. For two-outcome acts $[x \text{ on } A; y \text{ on } A^c]$ with $x \succsim y$, a constant-linear intrinsic integral has the following simple "biseparable" representation (Ghirardato-Marinacci (2001))

$$\widehat{\rho}\,(u \circ f) = u(x)\rho(A) + u(y)\,(1 - \rho(A))\,.$$

Turning to the third part of the characterization, it may seem a bit surprising that utility sophistication respectively the underlying Tradeoff Consistency axiom entails non-trivial restrictions on betting preferences, by ruling out betting preferences satisfying either P4 or both Union- and Splitting Invariance, but not both. To see how this is possible, note that while utility sophistication by iteslf does not restrict betting preferences *for given stakes* x and y, it does constrain betting preferences across stakes, even in the absence of P4. The existence of such restrictions explains how the imposition of *further* restrictions on betting preferences across stakes such as P4 can entail restrictions on betting preferences for given stakes.

The two Invariance axioms are intuitive and have intrinsic appeal even in the absence of utility sophistication. They can be viewed as capturing a notion of ambiguity attitudes as invariant across events with the same structure of ambiguity. This is perhaps brought out more clearly by the following implication (indeed, characterization) of Union- plus Splitting Invariance in the convex-ranged case.

For all $A, B \in \Sigma$ and $\alpha, \beta \in [0,1]$ such that $\pi(A) = \alpha\pi(B) + \beta$ for all $\pi \in \Pi$, $\rho(A) = \alpha\rho(B) + \beta$

In view of their appeal, it is not surprising that both conditions have some *in cameo* precedents in the literature. Union Invariance on the one hand is sufficiently intuitive for Epstein-Zhang (2001) to build it into their very definition of an event T as "revealed unambiguous" (cf. Nehring 2003).[8] Splitting Invariance as well is not entirely new, as it can be reformulated as a restriction on betting preferences over independent events. Say that events A and B are independent if $\pi(B/A) = \pi(B/A^c)$ for all $\pi \in \Pi$. If preferences are compatible with the convex-ranged context $\trianglerighteq$ as usual, then it can be shown easily that they satisfy Splitting Invariance if and only if

$$\rho(A \times B) = \rho(A)\rho(B)$$

[8] That is to say, Epstein-Zhang's definition of revealed unambiguous events is such that Union Invariance (applied to reveal unambiguous events instead of Λ) holds by definition.

for all $A \in \Sigma$ and $B \in \Lambda$ such that A and B are independent. In defining "product capacities" for independent events, authors such as Ghirardato (1997) and H. (#) have appealed to related "Fubini properties".[9]

5. UTILITY SOPHISTICATION IN PARTICULAR MODELS

5.1 Models with EU aggregators

Note that Utility Sophistication can be formulated essentially equivalently using the notion of an aggregegator Ψ of the expected-utility values under Π. Let $\mathbb{E}_\Pi : \mathcal{Z} \rightarrow [0,1]^\Pi$ denote the evaluation operator given by $\mathbb{E}_\Pi(Z) = (E_\pi Z)_{\pi \in \Pi}$ for $Z \in \mathcal{Z}$; slightly abusing notation, $\mathbb{E}_\Pi \circ u$ maps acts f to vectors of expected utilities $(E_\pi u \circ f)_{\pi \in \Pi}$. [10] Clearly, $\mathbb{E}_\Pi(\mathcal{Z})$ is a convex subset of $[0,1]^\Pi$. An EU aggregator is simply a monotone mapping $\Psi : \mathbb{E}_\Pi(\mathcal{Z}) \rightarrow [0,1]$. Then, up to minor technicalities, Utility Sophistication is equivalent to the existence of an EU aggregator Ψ, an operator $\mathbb{E}_\Pi$ and a utility-function u such that

$$f \succsim g \text{ if and only if } \Psi(\mathbb{E}_\Pi(u \circ f)) \geq \Psi(\mathbb{E}_\Pi(u \circ g)), \text{ for any } f, g \in \mathcal{F}.$$

A variety of models in the literature can be put in this form. For example, the Minimum Expected Utility ("MEU") model due to Gilboa-Schmeidler (1989) corresponds to the aggregator

$$\Psi(U) = \min_{\pi \in \Pi} U_\pi.$$

A natural generalization admitting ambiguity-seeking is derived from evaluating acts according to the entire range of expectations under Π, with

$$\Psi(U) = J\left(\min_{\pi \in \Pi} U_\pi, \max_{\pi \in \Pi} U_\pi\right),$$

for some monotone function J; see e.g. Jaffray (1989). Such "Interval EU" preferences satisfy P4 if and only if J is linear, i.e. if and only only if Ψ can be written as

$$\Psi(U) = \alpha \max_{\pi \in \Pi} U_\pi + (1-\alpha) \min_{\pi \in \Pi} U_\pi,$$

[9] Indeed, the Fubini properties in these contributions is substantially stronger as they applies to cases in which both A and B may be ambiguous.

[10] Note also that in view of the Krein-Milman theorem, one could restrict attention to the subvector $(E_\pi u \circ f)_{\pi \in Ext(\Pi)}$

Hurwicz's classical optimism-pessimism criterion.[11] Dubbed α-MEU, it has been axiomatized by Ghirardato et al. (2002) (2002) and Kopylov (2002).

Klibanoff et al. (2003) study preferences associated with aggregators

$$\Psi(U) = E_\mu \phi(U_\pi),$$

where μ is a probability-measure on Π, and $\phi : \mathbf{R} \to \mathbf{R}$ is strictly increasing and continuous (typically smooth). If ϕ is smooth, then bets $[x \text{ on } A; y \text{ on } A^c]$ with small stakes (i.e. such that $u(x) - u(y)$ is close to zero, are evaluated approximately according to $E_\mu \pi$. Thus, preferences will satisfy P4 if and only if ϕ is linear, i.e. SEU.

A fairly diverse class of preferences with an EU aggregator representation has been studied in Siniscalchi's (2003) "plausible priors" model; all of these satisfy P4 by construction.

Finally, somewhat outside the present framework in considering choice-functions are than weak orders, in Nehring's (1991,2000) "Simultaneous Expected Utility" model $\Psi(U)$ is the lexicographic minimum of appropriately renormalized expected utilities; the renormalization allows an interpretation of the solution as a bargaining solution among alternative selves associated with the extremal priors Π. All of the above contibutions are situated in variants of the Anscombe-Aumann (1963) framework. This is no accident: indeed, by translating these contributions into the present setting, we will see that the typical assumptions made imply that these models are utility-sophisticated with respect to the AA context $\trianglerighteq_{AA}$, indeed typically even standard utility-sophisticated.

5.2 Ambiguity in the Anscombe-Aumann Framework

The Anscombe-Aumann (1963) framework is distinguished by taking acts to be mappings from states to probability distributions of consequences, rather than simply as mappings from states to consequences as in the Savage (1954) framework. These probability distributions are interpreted as objective probabilities of the realizations of an external random device ("roulette lotteries") that is not part of the explicitly modeled state space. In section 2, we have restated this description as a convex-ranged context of explicit beliefs $\trianglerighteq_{AA}$. We will begin by showing how a preference relation over Savage acts can be rediscribed as a preference relation over Anscombe-Aumann (AA-) acts.

Formally, an AA-act F is a finite-valued Σ_1-measurable mapping from the subjective state space Ω_1 to the set of probability distributions on X with finite support $\mathcal{L}$. Let $\mathcal{F}^{AA}$ denote their set.

[11] This follows from Theorem 2 together with Ghirardato et al.'s (2002) demonstration of the equivalence of the c-linearity of I and the constant-α representation.

Denoting elements of $\mathcal{L}$ by $q = (q^x)_{x \in X}$, one can write $F = [q_1 \text{ on } S_1; q_2 \text{ on } S_2; ...]$ in analogy to the notation for Savage acts. Note that since Σ is the product algebra of Σ_1 and Σ_2, any Savage act f can be written in the form $[x_{i,j} \text{ on } S_i \times T_{i,j}]_{i \leq n, j \leq n_i}$. One can thus associate with any Savage act $f = [x_{ij} \text{ on } S_i \times T_{ij}]$ the AA-act $F = F(f) = [p_i \text{ on } S_i]$, with $p_i^x = \sum_{j \leq n_i, x_{ij} = x} \eta(T_{i,j})$. $F(f)$ is simply the AA-act that associates with any subjective state $\omega \in \Omega_1$ the lottery that yields the consequence x with unambiguous (subjective) probability entailed by the likelihood judgments $\trianglerighteq_{AA}$. By the convex-rangedness of $\trianglerighteq_{AA}$, this mapping is onto, i.e. any AA-act is the image of some Savage act.

In order to associate with the given preference relation $\succsim$ over Savage acts a well-defined preference relation over AA acts, one needs to extend the assumption that preferences are compatible with the context $\trianglerighteq_{AA}$ in the following natural way.

Axiom 10 (Strong Compatibility) *For all $f \in \mathcal{F}$, $x, y \in X$ and $A, B \subseteq C \in \Sigma$:*

i) $[x \text{ on } A; y \text{ on } C \backslash A;\ f \text{ elsewhere}] \succsim [x \text{ on } B; y \text{ on } B \backslash A;\ f \text{ elsewhere}]$ *if* $A \trianglerighteq B$ *and* $x \succsim y$, *and*

ii) $[x \text{ on } A; y \text{ on } C \backslash A;\ f \text{ elsewhere}] \succ [x \text{ on } A; y \text{ on } C \backslash A;\ f \text{ elsewhere}]$ *if* $A \triangleright B$ *and* $x \succ y$.

Note that Compatibility is simply Strong Compatibility restricted to the case of $C = \Omega$; in turn, Strong Compatibility is entailed by Utility Sophistication (with respect to the same context, of course).[12]

The lottery p *stochastically dominates* the lottery q if, for all $y \in X$, $\sum_{x: x \succsim_{AA} y} p^x \geq \sum_{x: x \succsim_{AA} y} q^x$; p *stochastically dominates* q *strictly* if at least one of these inequalities is strict. The AA-act $F = [p_i \text{ on } S_i]$ *(strictly) stochastically dominates* the AA-act $F = [q_i \text{ on } S_i]$ if p_i (strictly) stochastically dominates q_i for every i.

Fact 2 *The following two conditions are equivalent:*

i) $\succsim$ *is strongly compatible with* $\trianglerighteq_{AA}$

ii) For all f, g such that $F(f)$ stochastically dominates $F(g)$ (resp. strictly) $f \succsim g$ (resp. $f \succ g$).

It is immediate from part ii) that if $\succsim$ is strongly compatible with $\trianglerighteq_{AA}$, any f, f' such that $F(f) = F(f')$ must be indifferent. One thus obtains a well-defined weak order on $\mathcal{F}^{AA}$ by setting

$$F \succsim_{AA} G :\Leftrightarrow f \succsim g \text{ for any } f \text{ and } g \text{ such that } F = F(f) \text{ and } G = F(g).$$

[12] The second, strict part of Strong Compatibility plays no role in the following; we have included it only to ensure that Stong Compatibility entails Compatibility.

Furthermore, $\succsim_{AA}$ respects stochastic dominance. The following result is therefore a straightforward corollary of Fact 2; it implies that preferences over Savage acts that are strongly compatible with the context $\trianglerighteq_{AA}$ and preferences over AA-acts that respect AA Stochastic Dominance are essentially the same object.

Proposition 2 *If the preference ordering over Savage acts are strongly compatible with the context* $\trianglerighteq_{AA}$*, the associated preference ordering* $\succsim_{AA}$ *respects AA stochastic dominance. Conversely, if the preference ordering* $\overline{\succsim}$ *over AA-acts respects stochastic dominance, there exists a unique preference ordering* $\succsim$ *that is strongly compatible with* $\trianglerighteq_{AA}$ *such that* $\overline{\succsim} = \succsim_{AA}$ *.*

We will now show that the standard assumptions on AA preferences in contributions such as Schmeidler (1989) and Gilboa-Schmeidler (1989) amount to regular utility sophistication of the corresponding preferences over Savage acts. These assumptions are summarized by the following three axioms.

Axiom 11 (Monotonicity) *For all acts* $F \in \mathcal{F}^{AA}$*, lotteries* $p, q \in \mathcal{L}$ *and events* $S \in \Sigma_1$: $[p$ *on* $S; F(\omega)$ *elsewhere*$] \succsim_{AA} [q$ *on* $S; F(\omega)$ *elsewhere*$]$ *whenever* $x \succsim y$*.*

Axiom 12 (Lottery Independence) *For all lotteries* $p, q, r \in \mathcal{L}$ *and all* $\alpha \in (0, 1]$: $p \succsim_{AA} q$ *if and only if* $\alpha p + (1 - \alpha)r \succsim_{AA} \alpha q + (1 - \alpha)r$

Axiom 13 (Certainty Independence) *For all acts* $F, G \in \mathcal{F}^{AA}$*, constant acts (lotteries)* $H \in \mathcal{F}^{AA}_{const} = \mathcal{L}$ *and all* $\alpha \in (0, 1]$: $F \succsim_{AA} G$ *if and only if* $\alpha F + (1 - \alpha)H \succsim_{AA} \alpha G + (1 - \alpha)H$.[13]

The two main results of the paper, Theorem 1 and 2 yield the following result.

Proposition 3 *Suppose that the preference ordering* $\succsim$ *is bounded, solvable, Archimedean and strongly compatible with respect to the context* $\trianglerighteq_{AA}$*. Let* $\succsim_{AA}$ *denote the associate preference ordering over AA-acts. Then*

i) $\succsim$ *satisfies Tradeoff Consistency if and only if* $\succsim_{AA}$ *satisfies Monotonicity and Lottery Independence. Furthermore,*

ii) $\succsim$ *satisfies Tradeoff Consistency and Stake Independence (P4) if and only if* $\succsim_{AA}$ *satisfies Monotonicity and Certainty Independence.*

[13] As usual, this mixture-operation is defined point-wise

Proposition 3 yields a subjective, epistemic foundation of the standard modelling of ambiguity in the AA framework. All axioms are conditions on preferences over Savage acts, some of them conditioned on likelihood comparisons that are either given as an independent datum or imputed by the analyst. Since all uncertainty is treated on par, all (unconditioned) purely behavioral assumptions carry their usual, transparent meaning. By contrast, the original AA framework treats objective and subjective uncertainty differently; while the trick of including the objective uncertainty in the consequences is mathematically neat, it rather drastically changes the meaning of standard assumptions such as Monotonicity which turn out to be much stronger than usual.

This intransparency potentially affects assumptions made within this framework to characterize specific preference models. Epstein (1999), for example, criticizes Schmeidler's (1989) and Gilboa-Schmeidler's (1989) definition of ambiguity aversion as too restrictive and/or inapplicable in a Savage setting. In the companion paper Nehring (2004), we formulate a definition of ambiguity aversion in terms of betting preferences, and show that it yields Schmeidler's *in the utility-sophisticated case.* In the absence of utility-sophistication, for example in the context of the CEU model, it has however none of the restrictive and undesirable implications critized by Epstein.

Likewise, one obtains Schmeidler's (1989) "mixed CEU" model by imposing "Comonotonic Independence" restricted to non-random (Σ_1-measurable) acts; this simple observation throws light on the by know well-known fact that Schmeidler's (1989) "mixed CEU" model is quite distinct, indeed disjoint, from proper CEU models as formulated in a Savage framework (Gilboa 1987, Sarin-Wakker 1992).

The above epistemic sujective rendering of the AA setup is different from the recent preference-based translation by Ghirardato et al. (2003). The key to their work is a preference-based definition of utility-mixtures. It allows them to appeal to mathematically analogous axioms in a Savage context. However, since they now refer to different objects, namely Savage acts, the axioms are now different in content.

5.3 Choquet Expected Utility

A main contribution of Theorem 1 was to show that utility-sophistication with respect to a convex-ranged context implies bi-determinacy, i.e. the determination of preferences over general multi-valued acts from preferences over unambiguous acts (revealing a ranking of utility-differences) and preferences over bets. The Choquet Expected Utility (CEU) model in which acts are ranked according to the Choquet integral of utilities $\int u \circ f d\nu$ is the main alternative bi-determinate model in the literature; this bi-determinacy comes out especially clearly in Sarin-Wakker's (1992) axiomatization based on a Cumulative Dominance axiom which explicitly constructs multi-act preferences from preferences over bets. In contrast to utility sophistication, the CEU model is designed to allow for departures from expected utility also in the absence of ambiguity, accomodating for example the Allais (1953) paradox. If one writes the non-normalized capacity ν as $\phi \circ \rho$, this is expressed by the non-linearity of ϕ.

When are Choquet preferences utility-sophisticated? While this can happen when the underlying context is not convex-ranged (hence when utility-sophistaction does not entail bi-determinacy), this never happens under convex-rangedness in the presence of any ambiguity.

Proposition 4 *Suppose that a CEU preference ordering $\succsim$ is utility-sophisticated relative to the convex-ranged context $\trianglerighteq$; then $\succsim$ is in fact SEU.*

To illustrate the logic of the result, consider a utility-sophisticated CEU preference ordering represented by a non-symmetric capacity ρ, that is any capacity for which $\rho(A) + \rho(A^c) \neq 1$ for some set A; note that Ellsberg-style betting preferences imply specifically the inequality $\rho(A) + \rho(A^c) < 1$. The proof in the Appendix extends the following argument to general, possibly symmetric capacities. Denote consequences in (non-normalized) utiles, and take $B \subseteq A^c$ such that $B \equiv A^c \backslash B$. Let g denote the constant act $\mathbf{1}$, and f the act

$$[1 \text{ on } A, 2 \text{ on } B, 0 \text{ on } A^c \backslash B].$$

Conditional on A^c, this act entails an unambiguous 50-50 lottery with utility-payoffs 2 or 0. The act f has therefore expected utility of 1 for every prior $\pi \in \Pi$. By utility-sophistication preference, thus

$$f \sim g.$$

By EU maximization on unambiguous acts, $\nu = \rho$. Hence $\int u \circ f d\nu = 2\rho(B) + 1[\rho(A+B) - \rho(B)] = \rho(A+B) + \rho(B)$. Since CEU preferences satisfy P4, ρ must be affine invariant. From this one can deduce that in fact

$$\int u \circ f d\nu = 1 + \frac{1}{2}(\rho(A) + \rho(A^c) - 1). \tag{1}$$

Thus $f \sim g$ only if $\rho(A) + \rho(A^c) = 1$. Note that since the expression (1) is the certainty equivalent of the act f, its difference from the unambiguous expectation of 1 is therefore proportional to the extent of non-symmetry $|\rho(A) + \rho(A^c) - 1|$ and could be large (up to $\frac{1}{2}$).

Proposition 4 illustrates clearly that utility sophistication with respect to a convex-ranged context is much stronger than EU maximization over unambiguous acts. The Proposition implies in particular that a "mixed CEU" model a la Schmeidler is a proper CEU model only in the degenerate case when it is in fact SEU. This fact has been established before by Sarin-Wakker (1992), and could have been used to derive the Proposition instead of the direct proof given in the Appendix.

More broadly, Proposition 4 reveals a fundamental incompatibility between rank-dependence and utility-sophistication. It extends to non-P4 generalizations of CEU such as Cumulative Prospect Theory, and does not hinge on convex-rangedness of the context. If the context is not convex-ranged, there often exist some non-degenerate utility-sophisticated CEU preferences. But their set will in many cases still be fairly degenerate. For example, let μ_1 and μ_2 denote probability measures on two distinct subalgebras $\mathcal{A}_1, \mathcal{A}_2 \subseteq \Sigma$, and let $\mathcal{A}$ denote the smallest algebra containing both $\mathcal{A}_1$ and $\mathcal{A}_2$. Let Π^2 denote the set of priors π that agree with μ_i on $\mathcal{A}_i$. It follows from the analysis in Nehring (1999) that if a CEU preference relation is utility-sophisticated with respect to Π^2, the representing capacity must be additive on all of $\mathcal{A}$. Of course, utility-sophistication by itself carries no such implication, as evidenced by the MEU model; nor does the assumption of CEU preferences: in general, many CEU preferences are strongly compatible with $\trianglerighteq_{\Pi^2}$, without eliminating ambiguity about events in $\mathcal{A} \backslash (\mathcal{A}_1 \cup \mathcal{A}_2)$.

6. UTILITY-SOPHISTICATION AS A BEHAVIORAL PROPERTY

Utility Sophistication has been defined relative to an exogeneously specified set of likelihood comparisons $\trianglerighteq$. It is natural to ask whether utility sophistication can be understood in purely behavioral terms: (when) is it possible to ascertain on the basis of preferences alone that all departures from SEU maximization are due to ambiguity. We shall refer to such preferences as "revealed utility-sophisticated". If there was a canonical purely behavioral definition of "revealed probabilistic beliefs" as a coherent likelihood relation $\trianglerighteq^{\#}$, utility sophistication could in turn be defined behaviorally as utility sophistication relative to revealed beliefs $\trianglerighteq^{\#}$. While it seems doubtful that such a definition is possible in general, we shall show below that these difficulties can be frequently overcome; in particular, we will argue that there is a fairly strong case for considering considering any preference ordering as revealing utility-sophistication for which there exists *some* convex-ranged context relative to which it is utility sophisticated in the sense of section 3.

6.1. Occam's Razor and the Principle of Charity

Clearly, if the agents' true belief relation is not observed, there cannot be a "foolproof" way to infer it from behavioral data alone. For example, Karni (#) has argued that even in the context of Savage's SEU theory, Savage's definition does not necessarily identify the agents' true beliefs, since the agents true utility function may be . Roughly speaking, Karni argues that, for each state, only the product $\pi_\omega u_\omega(.)$ can be inferred from preferences; thus the agent's true probability π_ω agrees with the one identified by Savage's theorem only if the agents' true utility scale is state-independent.

While this is correct as it stands, the attribution of a state-independent utility-scale in the absence of any evidence to the contrary is far from arbitrary, and can be viewed as a natural application of Occam's razor. In the present context, it is natural to flesh out Occam's "principle of parsimony" razor as a "principle of charity", according to which observed behavior is to be explained as "intelligible", reason-guided as possible. Specifically, we will consider three criteria of reason-guidedness:

1. The precision of imputed beliefs. While greater precision is not a hall-mark of greater rationality per se, it does constrain preferences more tightly (via Strong Compatibility) and renders them thereby more "reason-guided".

2. The "sophistication" (tightness and "rationality") in the link between preferences and imputed beliefs. Here we take Strong Compatibility as the weakest form of sophistication, and Utility Sophistication as strongest form. Intermediate criteria such as respect for Union- and Splitting Invariance might conceivably be also invoked but will not be considered here.

3. If available, partial direct information about the DM's beliefs, or information about the probabilistic structure of the state-space, such as the use of an independent randomization device in the AA narrative.

An imputation/explanation exploiting such external information will be referred to as *semi-behavioral,* in contrast to a *fully behavioral* imputation that relies on preference information only. Throughout, we will assume a set of beliefs (coherent likelihood-relation) $\trianglerighteq$ to be *imputable* if the given preference relation is strongly compatible with them.

To begin with, consider the case of SEU preferences, with a convex-ranged "revealed likelihood" relation $\trianglerighteq^{rev}$ a la Savage. Here, any imputable set of beliefs $\trianglerighteq$ is contained in $\trianglerighteq^{rev}$, hence $\trianglerighteq^{rev}$ has no competition using the precision criterion alone. Moreover, preferences satisfy the highest standard of sophistication, utility sophistication, relative to these beliefs. Hence, application of the principle of charity explains these preferences unabmiguously as the result of utility-sophisticated choice relative to the precise probabilistic beliefs $\trianglerighteq^{rev}$.

By contrast, for probabilistically sophisticated non-SEU preferences, the principle of charity does not necessarily yield clearcut results since the first two criteria may conflict, as illustrated by the following example that has been discussed before in particular by Epstein-Zhang (2001) and Ghirardato-Marinacci (2002) .

Example 3. Let μ a convex-ranged probability measure, and $\phi : [0,1] \to [0,1]$ an increasing, strictly convex function mapping the unit interval onto itself. Suppose that preferences have a CEU representation with capacity $\phi \circ \mu$, and are thus probabilistically sophisticated in the sense of Machina-Schmeidler (1992). Thus, one natural explanation of these prefences is that the DM has probabilistic beliefs given by μ and is "probabilistically risk-averse". However, there is a competing explanation, namely that the decision-maker is utility-sophisticated but ambiguity-averse, evaluating acts according to the minimum expected utility of the core of the capacity ν, $\Pi = \{\pi : \pi \geq \nu\}$.[14] The first one is more "charitable" in attributing to the DM precise probabilistic beliefs. The second explanation, on the other hand, is more charitable in attributing to the DM a greater degree of sophistication in acting upon the probabilistic beliefs that it attributes to him. On the basis of preferences alone, there seems to be therefore no basis for priveleging one explanation over the other. At most, one could argue in favor of a "convention" by postulating the primacy of one criterion of intelligibility over the other, for example by declaring probabilistic sophistication to reveal absence of ambiguity by definition.[15]

We conclude from this example that the principle of charity can at most yield an unambiguous verdict on the revealed utility-sophistication of preferences for particular classes of preferences and/or in the presence of a sufficient amount of external information. We will begin with the second case as the easier one.

[14] The capacity ν is easily verified to be submodular; the existence of an alternative MEU represenation follows therefore from a result by Schmeidler (1989).

[15] This seems to be the line taken by Epstein-Zhang (2001). Ghirardato-Marinacci (2002), 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 non-behavioral information. They also point out how the existence of exogeneously identified unambiguous events allows one to distinguish between the two possible interpretations of a probabilistically sophisticated DM.

6.2. Revealed Utility Sophistication with External Information

Suppose, thus, that the DM's beliefs $\trianglerighteq$ are known to contain a convex-ranged context $\trianglerighteq^0$ such as the context $\trianglerighteq_{AA}$. The following Proposition, the fundamental result of this section, shows that if preferences are tradeoff consistent / utility sophisticated with respect to the context $\trianglerighteq^0$, then they are in fact tradeoff consistent / utility sophisticated with respect to the agents actual beliefs, whatever these are!

Proposition 5 *i) Suppose $\succsim$ is tradeoff consistent with respect to the convex-ranged context $\trianglerighteq^0$ and strongly compatible with $\trianglerighteq \supseteq \trianglerighteq^0$. Then $\succsim$ is tradeoff consistent with respect to $\trianglerighteq$.*

ii) Suppose that $\succsim$ is tradeoff consistent and Archimedean with respect to the convex-ranged context $\trianglerighteq^0$, bounded, solvable and strongly compatible with $\trianglerighteq \supseteq \trianglerighteq^0$. Then $\succsim$ is utility sophisticated with respect to $\trianglerighteq$.

Part ii) is a straightforward consequence of part i) and Theorem 1, in light of the fact that a preference relation is Archimeadean (resp. solvable) with respect to $\trianglerighteq^1$ whenever it is Archimedean (resp. solvable) with respect to $\trianglerighteq^0$. As demonstrated by Example 3 above, the convex-rangedness of $\trianglerighteq^0$ is crucial to the validity of the result. Proposition 5 motivates the following definition. Say that an agents' preferences $\succsim$ are **revealed utility sophisticated given** $\trianglerighteq^0$, if they are utility sophisticated relative to any imputable $\trianglerighteq \supseteq \trianglerighteq^0$. The conditions in part ii) of the Proposition provide an operational criterion of revealed utility sophistication.

Applied to the traditional Anscombe-Aumann setting, Propositions 3 and 5 imply that preferences are revealed utility sophisticated if and and only if the associated AA preferences are monotone. Thus a large number of existing models build in utility-sophistication.

6.3. Utility Sophistication Revealed by Preferences Alone

Let us return to the case in which only preference information is available to determine whether a decision maker's preferences are utitility sophisticated (that is, u.s. relative to plausibly imputable beliefs). Unless preferences are SEU, there will in general not exist a unique maximal imputable beliefs. Preferences will typically be utility sophisticated relative to some (maximal) imputable beliefs (say) but not relative to others that are not comparable by set-inclusion. In order to decide in favor of utility sophistication non-arbitrarily, one needs to ensure that there do not exist imputable beliefs that are "structurally more precise" in an appropriate sense, but relative to which preferences fail to be utility-sophisticated.

Given a criterion of greater stuructural precision that goes beyond mere set-inclusion, one can define a preference relation $\succsim$ as **strongly** (respectively **weakly**) **revealed utility sophisticated** if it is utility sophisticated with respect to some likelihood ordering $\trianglerighteq^1$ and "structurally more precise" (respectively structurally no less precise) than any other imputable likelihood ordering $\trianglerighteq^2$ relative to which $\succsim$ is not utility sophisticated.

It remains to flesh out the notion of structural precision. Such a notion would have to naturally exploit structural completeness features of the likelihood orderings compared. An obvious candidate for a relevant structural feature is the convex-rangedness the likelihood ordering. Thus, without further discussion, we postulate that any convex-ranged likelihood ordering is *structurally more precise* than any non-convex-ranged likelihood ordering. Two further natural structural features are the completeness of $\trianglerighteq$ (singletonhood of $\Pi_{\trianglerighteq}$) and the "finite-dimensionality" of $\trianglerighteq$ (that is: the containment of $\Pi_{\trianglerighteq}$ in the convex hull of a finite number priors). An important example of a finite-dimensional context is the context $\trianglerighteq_{AA}$ when the number of "generic states" is finite.

For the sequel, we will fix the structural precision ordering $\Supset$ over the coherent likelihood relations on the σ-algebra Σ as follows. To ensure that the following definition yields a genuine partial ordering, we shall restrict the domain of $\Supset$ to relations $\trianglerighteq$ such that $\Pi_{\trianglerighteq}$ contains a dense set of countably additive priors; we shall refer to such relations as "*σ-additive*".[16]

Definition 2 (Structural Precision) $\trianglerighteq^1 \Supset \trianglerighteq^2$ *if* $\Pi_{\trianglerighteq^1}$ *is a singleton and* $\Pi_{\trianglerighteq^2}$ *is not, or if* $\Pi_{\trianglerighteq^1}$ *is finite-dimensional and* $\Pi_{\trianglerighteq^2}$ *is not, or if* $\Pi_{\trianglerighteq^1}$ *is convex-ranged and* $\Pi_{\trianglerighteq^2}$ *is not, or if* $\trianglerighteq^1 \supset \trianglerighteq^2$.[17]

We have the following fully behavioral counterpart to Proposition 5.

Proposition 6 *Suppose* $\succsim$ *is strongly compatible with the convex-ranged σ-additive contexts* $\trianglerighteq^1$ *and* $\trianglerighteq^2$ *one of which is finite-dimensional, and such that all* $\succsim$*-null events are* $\trianglerighteq^1$*- and* $\trianglerighteq^2$*-null. If it is furthermore bounded, solvable, Archimedean and tradeoff-consistent with respect to* $\trianglerighteq^1$ or $\trianglerighteq^2$*, then it is utility sophisticated with respect to* both $\trianglerighteq^1$ *and* $\trianglerighteq^2$.

[16] Just as in the complete case, $\sigma-$additivity is a largely technical requirement. It is much weaker than requiring all priors in $\Pi_{\trianglerighteq}$ to be countably additive, rather than merely a dense set; see Marinacci et al. (2003). Note however that if $\Pi_{\trianglerighteq}$ has a finite number of extreme points, both requirements are equivalent.

[17] This definition is meant to guide the following discussion and not written in stone. In particular, it would probably be affected by the consideration of almost-convex-ranged contexts.

The assumption that all $\succsim$-null events are $\trianglerighteq^i$-null means that all practical certainties are underwritten as epistemic certainties. This assumption is without essential loss of generality, in that one can show that, under appropriate continuity assumptions on preferences, if a convex-ranged context $\trianglerighteq$ is strongly compatible with $\succsim$, there exists a super-context $\trianglerighteq' \supseteq \trianglerighteq$ for which all $\succsim$-null events are $\trianglerighteq'$-null and that is also strongly compatible with $\succsim$.

In view of Proposition 5, Proposition 6 is a straighforward consequence of the following Fact and the observation that Solvability and Archimedeanicity are inherited by convex-ranged subcontexts.

Fact 3 *If $\trianglerighteq^1$ and $\trianglerighteq^2$ are convex-ranged σ-additive contexts on the σ-algebra Σ with identical null-sets and such that one of them finite-dimensional, then $\trianglerighteq^1 \cap \trianglerighteq^2$ is convex-ranged.*

Fact 3 is an application of Lyapunov's Theorem. It would fail to be true without the assumption that at least one context is finite-dimensional[18]. On the other hand, we do not know whether the same holds for Proposition 6.

Using the above terminology and definition of $\Supset$, Proposition 6 implies the following corollary (fudging the non-issue of $\sigma-$additivity).

Corollary 1 *Suppose that $\succsim$ is Archimedean and utility sophisticated relative to some $\sigma-$additive, convex-ranged (respectively finite-dimensional) context $\trianglerighteq$, with an onto utility function $u : X \to [0,1]$. Then $\succsim$ is in weakly (respectively strongly) revealed utility sophisticated.*[19]

To take stock, consider again the AA framework with a finite number of generic states and a σ-additive measure μ over random events. Suppose that we have verified that preferences satisfy all the maintained regularity assumptions plus tradeoff consistency with respect to $\trianglerighteq^1$; suppose also that we find that preferences are strongly compatible but not utility-sophisticated with respect to some other context $\trianglerighteq^2$. Then, by Propostion 6, $\trianglerighteq^2$ cannot be convex-ranged. Imputing this context $\trianglerighteq^2$ rather than $\trianglerighteq_{AA}$ (and possibly other beliefs) would be *uncharitable* in terms of both criteria of charity, precision and sophistication. There is thus a clearcut case for attributing the beliefs $\trianglerighteq_{AA}$ and utility sophistication to the DM on the basis of preference information alone.

Proposition 6 makes a weaker and more tentative case for attributing revealed utility-sophistication based if the context relative to which preferences are utility-sophisticated is merely convex-ranged. The non-arbitrariness of this attribution would be seriously cast into doubt in cases in which there

[18]Note that this confirms that finite-dimensionality is indeed a distinct structural property

[19]Weakly (respectively strongly) revealed utility sophistication are understood here as focussed on σ-additive contexts only.

existed an alternative imputable context $\trianglerighteq^2$ that is structurally more precise than $\trianglerighteq^1$, and with respect to which preferences fail to be utility-sophisticated. In view of the Proposition, this would presuppose the existence of a precision that is intermediate between convex-rangedness and finite-dimensionality, yet weak enough not to imply the conclusion of Proposition 6. Obviously, whether this is the case will need to be borne out by future investigations. Based on results in Nehring (2001), we suspect, however, that in many cases, there will exist a unique maximal convex-ranged context with which a preference relation is strongly compatible; in all those cases, preferences would be stongly revealed utility sophisticated.

The existence-clause in Corollary 1 raises the question how one can determine for which preferences it is satisfied. We will now provide an operational criterion for this in the case of stake-independent preferences satisfying.

If preferences are P4, then they must be c-linear in order to be utility-sophisticated relative to a convex-ranged context. If preferences are indeed c-linear, the existence of such a context can be determined via the notion of a maximal independent subrelation $\succsim^*$ of the given preference relation $\succsim$. Define the following mixture-operation $\alpha f \oplus (1-\alpha) g$ on the space of acts: for $\alpha \in [0,1]$, $\alpha f \oplus (1-\alpha) g$ denotes any act h such that, for all $\omega \in \Omega$, $u(h_\omega) = \alpha u(f_\omega) + (1-\alpha) u(g_\omega)$. Since utility-functions in a c-linear representation are unique up to positive affine transformation, the mixture-operation is well-defined in terms of preferences; Ghirardato et al. (2003) provide a directly behavioral definition of this operation. Also, in view of Eventwise Monotonicity, the choice of the act h is immaterial. A (possibly incomplete) relation $\succsim'$ is independent if, for all f, g, h and $\alpha \in (0,1]$, $f \succsim' g$ if and only if $\alpha f \oplus (1-\alpha) h \succsim' \alpha g \oplus (1-\alpha) h$. In Nehring (2001), we have obtained (a variant of) the following result.[20]

Proposition 7 (Nehring 2001) *Suppose that the preference ordering $\succsim$ has a c-linear representation $I \circ u$ such that $u(X)$ is convex.*

i) Then there exists a unique maximal independent subrelation $\succsim^*$, with

$$f \succsim^* g \text{ if and only for all } h \text{ and all } \alpha \in (0,1],\ \alpha f \oplus (1-\alpha) h \succsim \alpha g \oplus (1-\alpha) h.$$

ii) There exists a unique closed, convex set of priors Π^* such that

$$f \succsim^* g \text{ if and only } E_\pi u \circ f \geq E_\pi u \circ g \text{ for all } \pi \in \Pi^*.$$

[20] A first version of this result was presented in the talk Nehring (1996); the exact version of the characterization of $\succsim^*$ in i) was arrived at independently by Ghirardato et al. (2003).

In particular, Π^* *is the unique minimal set closed, convex of priors* Π *such that* $\succsim$ *is utility-sophisticated with respect to* Π *and* u*.*

Note the reference to the utility-function u inherited from the c-linear representation; that this u is indeed the "canonical" one can be justified by appealing to Ghirardato et al.'s (2003) behavioral definition of the mixture-operation $\oplus$, for u is uniquely identified (up to positive affine transformations, of course) by the condition that $z = \frac{1}{2}x \oplus \frac{1}{2}y$ implies $u(z) = \frac{1}{2}u(x) + \frac{1}{2}u(z)$.

Associate with $\succsim^*$ the coherent likelihood relation given via the restriction of $\succsim^*$ to bets. That is, define

$$A \trianglerighteq^* B \text{ if } [x \text{ on } A, y \text{ on } A^c] \succsim^* [x \text{ on } B, y \text{ on } B^c] \text{ for any } x, y \text{ such that } x \succ y.$$

Note that $\trianglerighteq^* = \trianglerighteq_{(\Pi^*)}$. Since for any $\trianglerighteq \not\subseteq \trianglerighteq^*, \Pi_{\trianglerighteq} \not\supseteq \Pi^*$, $\trianglerighteq^*$ is the unique maximal coherent likelihood relation $\trianglerighteq$ such that $\succsim$ is utility-sophisticated with respect to $\trianglerighteq$ and u. In genral,$\Pi_{(\trianglerighteq^*)} \supseteq \Pi^*$, and it may easily be the case that the inclusion is strict, since $\trianglerighteq^*$ typically contains less information than $\succsim^*$. However, if Π^* is convex-ranged, then by the uniqueness result in Nehring (2003, Theorem 2) one $\Pi_{(\trianglerighteq^*)} = \Pi^*$; it follows that Π^* is convex-ranged if and only if $\trianglerighteq^*$ is. Thus, Proposition 7 entails the following operational characterization of revealed utility-sophistication.

Proposition 8 *Suppose that the preference ordering* $\succsim$ *has a c-linear representation* $I \circ u$ *such that* $u(X)$ *is convex. Then* $\succsim$ *is utility-sophisticated with respect to some convex-ranged context* $\trianglerighteq$ *if and only if* $\trianglerighteq^*$ *is convex-ranged.*

Note that, in contrast to Proposition 5 above, "revealed utility-sophistication" does not assume u.s. relative to u. This is possible, since $\succsim$ is utility-sophisticated with respect to some u' and convex-ranged $\trianglerighteq$, preferences must have a c-linear representation $I' \circ u'$; by the cardinal uniquess of these representations, u' must be a positive affine transformation of u. Recall also that if Σ is a σ-algebra, $\trianglerighteq^*$ is convex-ranged if, for each $A \in \Sigma$ there exists $B \subseteq A$ such that $B \equiv^* A \backslash B$.

6.4. A Two-Classification of Preference Models

The above discussion allows to classify a variety of models proposed in the literature according to their utility-sophistication and stake-independence. To tie in directly with the literature, we assume the epistemized version of the traditional AA model analyzed in section 5.2. using Proposition 3, recalling that in view of Proposition 5, utility-sophistication can be identified with with respect to the context . Note that on this understanding, the MEU model due to Gilboa-Schmeidler (1989) is not the class of MEU preferences over Savage acts (as axiomatized by Casadesus et al. (2000) and Ghirardato et al. (2003)) that are strongly compatible with $\trianglerighteq_{AA}$, as shown by example 1. As noted in section 5.3, only CEU preferences (and rank-dependent more generally) can be compatible with the $\trianglerighteq_{AA}$ as a convex-ranged context only when they are degenerate (SEU). We have discussed various non-P4 and their utility-sophistication above. In particular, we noted that utility-smooth preferences will in general violate P4. By contrast, preferences (i.e. preferences that are locally linear in events) as in the work of Machina (2002) and Epstein (1999) do not conflict with P4 per se; for example, event-smoothness imposes only mild conditions on CEU preferences. However, in view of the identity of utility-evaluation functional I and the implicit integral $\widehat{\rho}$ under utility-sophistication, event-smoothness and utility-smoothness are essentially the same thing in this case, which implies the absence of non-degenerate event-smooth models that are utility-sophisticated as well as stake-independent.[21] Finally, while stake-independence is generally lost under conditioning, utility-sophistication is preserved (with respect to the original context or the one updated via Generalized Bayes' rule).

	–	**Utility Sophistication**
–	Cumulative Prospect Theory (KT 92) Utility-Smooth Preferences Event-Smooth Preferences (Machina 02) Conditional Pref. (Epstein-Le Breton 93)	Interval Expected Utility (Jaffray 89) Utility-Smooth Preferences (KMM 02) Event-Smooth Preferences (Machina 02) Conditional Util. Soph. Preferences
P4	CEU (Gilboa 89, Sarin-Wakker 92) Event-Smooth Preferences (Machina 02)	CEU (Schmeidler 89) MEU (Gilboa-Schmeidler 89) α-MEU (GMM 02, Kopylov 02) Plausible Priors (Siniscalchi 03) Simultaneous EU (Nehring 91, 00)
Probabilistic Soph.	Machina-Schmeidler (92)	Subjective EU (Savage 54)

Table 1: Two-Way Classification of Preference Models

[21] This claim would appear to be robust to the particular formalization of event-smoothness adopted. Epstein (#) and Machina (2002) define event-smoothness relative to an additive reference measure; in the present setting, it would be natural to use the sub-additive upper-probability π^{+}_{AA} for this purpose.

7. CONCLUSION

Rather than summarizing the paper, we conclude by mentioning three directions of future research in which the present paper could be taken further. First, throughout the paper, convex-ranged belief contexts have played a central role, in three distinct ways: they serve to derive utility-sophistication from trade-off consistency, they ensure the determination of multi-act preferences from preferences over binary acts (bets and 50-50 "lotteries"), and they allow to conceptualize utility-sophistication in purely behavioral terms. In each instance, the power of convex-rangedness comes from the entailed mixture-space structure over events and acts. Without that structure, none of the main results of this paper could have been obtained. However, as shown in Nehring (2003) inspired by Machina (2001), convex-rangedness is not strictly necessary, since this mixture-space structure already obtains when beliefs are merely almost-convex-ranged, which in turn seems to be by and large necessary for that structure. It is therefore a natural task to extend the results of this paper to almost-convex ranged contexts. While we do not anticipate deep issues, the task is not merely technical, since many axioms and characterizations will need to be generalized to accommodate the absence of unambiguous events.

A second natural follow-up challenge is the development of specific utility-sophisticated models. Effectively, in view of the above classification of models in the literature, quite a bit has effectively already been done. However, the bideterminacy implied by utility sophistication on convex-ranged contexts promises improvements in the conceptual foundations of existing models (and perhaps suggest new ones), in that it implies that all specific assumptions on utility-sophisticated preferences boil down to assumptions on betting preferences. This is illustrated by Theorem 2 of this paper, which characterizes a familiar restriction on multi-act preferences (c-linearity of the representation) in terms of two alternative conditions on betting preferences. Specification via assumptions on betting preferences has the substantial conceptual advantage of making the additional assumptions "orthogonal" to the assumption of utility-sophistication itself. For example, the multi-act formulation of ambiguity aversion in Gilboa-Schmeidler's (1989) axiomatization of the MEU model originally due to Schmeidler (1989) has lead to some controversy as to whether Schmeidler's definition may too strong to be generally applicable, especially in view of its tensions with the CEU model (see E-K?, Epstein (1999), GM (2001), and Klibanoff). However, restricted to betting preferences, the MEU model amounts to requiring that bets are evaluated according to their lower probability, i.e. $\rho(A) = \min_{\pi \in \Pi'} \pi(A)$. This restriction can be axiomatized naturally using a notion of ambiguity aversion defined in terms of betting preferences only (exploiting the existence of convex-ranged reference context); see Nehring (2001). It is clearly orthogonal to the assumption of utility-sophistication itself; for instance, it can be combined equally well with the CEU model.

Finally, utility-sophistication itself may possibly be generalized in an interesting manner, as it can be viewed as combining a Pareto-criterion across the evaluations based on different admissible priors with the assumption that the DM evaluates unambiguous acts according to expected utility. For descriptive purposes, it would seem naturally attractive to maintain the former but abandon the latter, in order to accomodate all manner of probabilistic risk-attitudes. However, since in these case the natural separability properties of expected utility would be lost, there is no obvious counterpart to tradeoff consistency in the general case; this may in fact be more than a mathematical difficulty, and may mean that the Pareto-criterion across admissible priors is less compelling in the presence of probabilistic risk-attitudes.

APPENDIX: PROOFS

Proof of Fact 1.

Sufficiency is obvious. To verify necessity, there exist events $A^{'} \subseteq A$, $B^{'} \subseteq B$ and $B'' \subseteq B^c$ such that $A^{'} \equiv A \backslash A^{'}$, $B^{'} \equiv B \backslash B^{'}$, and $B^{'''} \equiv B^c \backslash B''$. By construction, $B^{'} + B'' \in \Lambda$, with $\overline{\pi}\left(B^{'} + B''\right) = \frac{1}{2}$. Hence by the convex-rangedness of $\overline{\pi}$ on Λ, there exists T disjoint from $B^{'} + B''$ such that $0 < \overline{\pi}(T) \leq \frac{\min_{\pi \in \Pi} \pi(A) - \pi(B)}{2}$; note that the latter is strictly positive by the weak compactness of Π. By construction, for any $\pi \in \Pi$, $\pi\left(A^{'}\right) \geq \pi\left(B^{'}\right) + \pi(T)$, as needed to be shown. $\square$

Proof of Proposition 1. That ii) implies i) is straightforward; as to the Archimedean property, merely note that $I-$continuity implies an analogous property for decreasing sequences $\{A_n\}$ by switching the roles of x and y.

For the converse, take any $g \in \mathcal{F}$. By Eventwise Monotonicity and boundedness, $x^- \precsim g \precsim x^+$. Hence from the Archimedeanicity of $\succsim$ applied to the case $A + B = \Omega$, one can infer the existence of an event $T_g \in \Lambda$ such that $g \sim [x^+, T_g; x^-, T_g^c]$. By compatibility, all such events T_g have the same unambiguous probability $\overline{\pi}(T_g)$. For any consequence/constant act z, set $u(z) := \overline{\pi}(T_z)$. The mapping $V : g \to \overline{\pi}(T_g)$ represents $\succsim$ by construction. By Eventwise Monotonicity, V can be written as $I \circ u$, with I monotone and compatible with $\trianglerighteq$; note that I is normalized by construction; moreover, the uniqueness claim is straightforward from Solvability which implies that u is onto. It remains to verify that I is event-continuous.

To see this, consider $\{A_n \in \Lambda_E\}$ and $A \in \Lambda_E$ such that $\overline{\pi}(A_n/E)$ converges to $\overline{\pi}(A_n/E)$ and such that the family is $\{A_n\} \cup A$ is ordered by set-inclusion. Take any x, y, Z. W.l.o.g. $x \geq y$. It clearly suffices to show convergence of $I(x1_{A_n} + y1_{E \backslash A_n} + Z1_{E^c})$ for the case of $\{A_n\}$ being an increasing or decreasing sequence. The proof for both cases is analogous; assume the former, and that. Suppose that the claim is false, i.e., in view of the monotonicity of I that $\sup_{n \in N} I(x1_{A_n} + y1_{E \backslash A_n} + Z1_{E^c}) < I(x1_A + y1_{E \backslash A} + Z1_{E^c})$. By normalization, there exist an event $T \in \Lambda$ such that $\sup_{n \in N} I(x1_{A_n} + y1_{E \backslash A_n} + Z1_{E^c}) < I(1_T) < I(x1_A + y1_{E \backslash A} + Z1_{E^c})$. Hence, by Archimedeanicity, there exist $A' \in \Lambda_E$ and $A' \triangleleft^0 A$ such that $I(1_T) < I(x1_{A'} + y1_{E \backslash A'} + Z1_{E^c})$. But by the convergence assumption, $A' \trianglelefteq^0 A_n$ for some n, hence $I(x1_{A'} + y1_{E \backslash A'} + Z1_{E^c}) \leq \sup_{n \in N} I(x1_{A_n} + y1_{E \backslash A_n} + Z1_{E^c}) < I(1_T)$, a contradiction. $\square$

In the following Lemma, we state a key mathematical property of the intrinsic integral $\widehat{\rho}$ that will be used repeatedly in the sequel. Say that $Z \in \mathcal{Z}$ is $\trianglerighteq$-unambiguous conditional on the finite partition $\mathcal{S}$ if, for all $S_i \in \mathcal{S}$, $Z1_{S_i}$ is Λ_{S_i}-measurable; let $\mathcal{Z}_{\mathcal{S}}$ denote their class. For $Z \in \mathcal{Z}_{\mathcal{S}}$, the expectation conditional on $\mathcal{S}$ $E^0(Z/\mathcal{S})$ is any random variable ζ such that

$$\zeta(\omega) = \sum_{z\in[\mathbf{0},\mathbf{1}]} z\pi^0(\{\omega' \in S_i \mid Z(\omega') = z\}/S_i)(\omega) \text{ if } \omega \in S_i \text{ for some non-null } S_i;$$

let their set be denoted by $E^0(Z/\mathcal{S})$.

Lemma 1 *$\widehat{\rho}$ is the unique mapping $\widetilde{\rho} : \mathcal{Z} \to [0,1]$ such that*

i) For any event $A \in \Sigma$, $\widetilde{\rho}(1_A) = \rho(A)$, and

ii) **(Conditional Linearity)** *For any partition $\mathcal{S}$ and any $Z \in \mathcal{Z}_{\mathcal{S}}$, $\widetilde{\rho}(Z) = \widetilde{\rho}(\zeta)$ for any $\zeta \in E^0(Z/\mathcal{S})$.*

Note that Conditional Linearity implies in particular that $\widehat{\rho}$ restricted to unambiguous random variables is the ordinary expectation with respect to $\overline{\pi}$ or equivalently ρ.

<u>Proof of Lemma 1.</u>

To verify the Conditional Linearity of $\widehat{\rho}$, write Z as $\sum_{i,j} z_{ij}1_{A_{ij}}$ with $S_i = \sum_{j\leq n_j} A_{ij}$ 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 non-null S_i and all $\pi \in \Pi$,

$$\pi(C \cap S_i) = \sum_j \pi(C \cap A_{ij}) = \sum_j z_{ij}\left(\overline{\pi}(A_{ij}/S_i)\,\pi(S_i)\right) = \left(\sum_j z_{ij}\overline{\pi}(A_{ij}/S_i)\right)\pi(S_i),$$

which implies that $C \in [\zeta]$ for any $\zeta \in E^0(Z/\mathcal{S})$. Thus indeed $C \in [Z] \cap [\zeta]$, and therefore

$$\widehat{\rho}(Z) = \rho(C) = \widehat{\rho}(\zeta).$$

Conversely, assume that $\widetilde{\rho}$ satisfies i) and ii). Consider any $Z = \sum_i z_i 1_{S_i}$ and any $C \in [Z]$. By definition of $[Z]$, $1_C \in \mathcal{Z}_{\mathcal{S}}$ with $Z \in E^0(1_C/\mathcal{S})$. Hence

$$\widetilde{\rho}(Z) = \widetilde{\rho}(1_C) = \rho(C),$$

which establishes that $\widetilde{\rho} = \widehat{\rho}$. $\square$

<u>Proof of Theorem 1.</u>

<u>iii) implies ii)</u> To show that $\succsim$ is utility sophisticated with respect to $\trianglerighteq$, take any f, g such that $E_\pi u \circ f \geq E_\pi u \circ g$ for all $\pi \in \Pi$, and take $A \in [u \circ f]$ and $B \in [u \circ g]$. By construction, $\pi(A) \geq \pi(B)$ for all $\pi \in \Pi$, and therefore by the compatibility of ρ

$$\widehat{\rho}(u \circ f) = \rho(A) \geq \rho(B) = \widehat{\rho}(u \circ g),$$

i.e. $f \succsim g$. By the same token, if $E_\pi u \circ f > E_\pi u \circ g$ for all $\pi \in \Pi$, then $f \succ g$.

To verify that $\succsim$ is Archimedean, in view of Proposition 1 we need to verify that $\widehat{\rho}$ is event-continuous. Thus, take some $x, y \in X$, $Z \in \mathcal{Z}$, $E \in \Sigma$, $A \in \Lambda_E$ and some increasing sequence $\{A_n\}$ of events contained in A such that $\overline{\pi}(A_n/E)$ converges to $\overline{\pi}(A/E)$; we need to show that $\widehat{\rho}(x1_{A_n} + y1_{E\setminus A_n} + Z1_{E^c})$ converges to $\widehat{\rho}(x1_A + y1_{E\setminus A} + Z1_{E^c})$. By conditional linearity (Lemma 1),

$$\widehat{\rho}(x1_A + y1_{E\setminus A} + Z1_{E^c}) = \widehat{\rho}((\overline{\pi}(A/E)\,x + (1 - \overline{\pi}(A/E))\,y)\,1_E + Z1_{E^c})$$

and likewise

$$\widehat{\rho}(x1_{A_n} + y1_{E\setminus A_n} + Z1_{E^c}) = \widehat{\rho}((\overline{\pi}(A_n/E)\,x + (1 - \overline{\pi}(A_n/E))\,y)\,1_E + Z1_{E^c}).$$

Suppose $x \geq y$; the converse case is dealt with symmetrically.

Take $B \in [Z1_{E^c}]$, $A' \in [(\overline{\pi}(A/E)\,x + (1 - \overline{\pi}(A/E))\,y)\,1_E]$ and an increasing sequence of events $A'_n \in [(\overline{\pi}(A_n/E)\,x + (1 - \overline{\pi}(A_n/E))\,y)\,1_E]$ contained in A'. By construction,

$$\widehat{\rho}((\overline{\pi}(A/E)\,x + (1 - \overline{\pi}(A/E))\,y)\,1_E + Z1_{E^c}) = \rho\,(A' + B)\,,$$

as well as

$$\widehat{\rho}((\overline{\pi}(A_n/E)\,x + (1 - \overline{\pi}(A_n/E))\,y)\,1_E + Z1_{E^c}) = \rho\,(A'_n + B)$$

Note that by definition, A' and the A'_n are unambiguous given E. Hence by the event-continuity of ρ, $\rho\,(A' + B) = \lim_{n\to\infty} \rho\,(A'_n + B)$, and therefore

$$\widehat{\rho}(x1_A + y1_{E\setminus A} + Z1_{E^c}) = \lim_{n\to\infty} \widehat{\rho}(x1_{A_n} + y1_{E\setminus A_n} + Z1_{E^c}),$$

as needed to be shown.

ii) implies i) It is clear the Utility Sophistication implies Compatibility. To verfiy Tradeoff Consistency, take any $x, y, x', y' \in X, f, g \in \mathcal{F}$ and events A disjoint from B and A' disjoint from B' such that $A \equiv B \rhd \emptyset$ and $A' \equiv B'$ and such that [x on $A; y$ on $B; f(\omega)$ elsewhere] $\succsim$ [x' on $A; y'$ on $B; f(\omega)$ elsewhere]. By the assumption on A and B, for all $\pi \in \Pi$, $\pi\,(A) = \pi\,(B) > 0$; therefore, if it was the case that $u(x) + u(y) < u(x') + u(y')$, then the strict part of Utility Sophistication would imply that [x on $A; y$ on $B; f(\omega)$ elsewhere] $\prec$ [x' on $A; y'$ on $B; f(\omega)$ elsewhere], which is false by assumption. Thus $u(x) + u(y) \geq u(x') + u(y')$, which implies by the non-strict part of Utility Sophistication that [x on $A'; y$ on $B'; g(\omega)$ elsewhere] $\succsim$ [x' on $A'; y'$ on $B'; g(\omega)$ elsewhere], as needed to be shown.

i) implies iii)

Since Tradeoff Consistency implies Eventwise Monotonicity, by Proposition 1 there exist an onto function $u : X \to [0, 1]$ and a normalized functional $I : \mathcal{Z} \to [0, 1]$ that is monotone, event-continuous and compatible with $\trianglerighteq$ such that $f \succsim g$ if and only if $I(u \circ f) \geq I(u \circ g)$, for all $f, g \in \mathcal{F}$. In particular, ρ is event-continuous as the restriction of I to indicator-functions. It remains to show that $I = \widehat{\rho}$.

Step 1. 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 2 *For any $\alpha \in \mathbf{D}, w, x, y \in X, A, B, T \in \Sigma$ such that $\overline{\pi}(T) = \overline{\pi}(A/B) = \alpha$: 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 $\alpha \geq \frac{1}{2}$; 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 $T_1 + T_2 + T_3 = \Omega, T_2 + T_3 = T$, and $\overline{\pi}(T_2) = \frac{1}{2}\beta$, hence also $\overline{\pi}(T_3) = \frac{1}{2}$ and $\overline{\pi}(T_2/T_1 + T_2) = \beta$. Likewise, define events A_1, A_2, A_3 such that $A_1 + A_2 + A_3 = B, A_2 + A_3 = A$, and $\overline{\pi}(A_2/B) = \frac{1}{2}\beta$, hence also $\overline{\pi}(A_3/B) = \frac{1}{2}$ and $\overline{\pi}(A_2/A_1 + A_2) = \beta$. Such events exist by the convex-rangedness of $\trianglerighteq$.

Take any $D \in \Lambda$ such that $\overline{\pi}(D) = \beta$, 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

$$[z, T_1 + T_2; x, T_3] \sim [w, T_1 + T_2; w, T_3]. \tag{2}$$

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 (2). By transitivity, we get

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

as desired.

Step 2. We shall next obtain the desired conclusion for the subset dyadic-valued functions $Y \in \mathcal{Z}$, which we shall abbreviate to $\mathcal{Z}_{\mathbf{D}}$. Thus, take any $Y = \sum_{i \leq n} y_i 1_{E_i} \in \mathcal{Z}_{\mathbf{D}}$; by solvability, there exists $f = [w_i, E_i]_{i\leq n} \in \mathcal{F}$ such that $u(w_i) = y_i$ for all i, so that $Y = u \circ f$. For each $i \leq n$, pick $A_i \subseteq E_i$ such that $\overline{\pi}(A_i/E_i) = u(w_i)$. By $n-$fold application of Lemma 2, $f \sim \left[x^+, \sum_{i\leq n} A_i; x^-, \left(\sum_{i\leq n} A_i\right)^c\right]_{i\leq n}$. Since $\sum_{i\leq n} A_i \in [Y]$ by construction, one obtains

$$I(Y) = I(u \circ f) = \rho(\sum_{i\leq n} A_i) = \widehat{\rho}(Y),$$

demonstrating that $I = \widehat{\rho}$ on $\mathcal{Z}_{\mathbf{D}}$.

Step 3.

This conclusion is extended to all of $\mathcal{Z}$ by an inductive continuity argument. Let $\mathcal{Z}_k$ the set of random variables $Y \in \mathcal{Z}$ such that in their canonical representation $Y = \sum_{i\leq n} y_i 1_{E_i}$ no more than k y_i's are not dyadic. Step 2 has established that $I = \widehat{\rho}$ on $\mathcal{Z}_{\mathbf{D}} = \mathcal{Z}_0$. Suppose therefore that $I = \widehat{\rho}$ on $\mathcal{Z}_k$; we need to show that $I = \widehat{\rho}$ on $\mathcal{Z}_{k+1}$. Take $Y = \sum_{i\leq n} y_i 1_{E_i} \in \mathcal{Z}_{k+1}$, and assume w.l.o.g. that $y_1 \in (0, 1]\backslash\mathbf{D}$.

Take an increasing sequence $\{v_j\}$ in $\mathbf{D}$ converging to y_1, and take $B \in \left[\sum_{2\leq i\leq n} y_i 1_{E_i}\right]$, $A \in [y_1 1_{E_1}]$ and an increasing sequence $\{A_j\}$ contained in A such that $A_j \in [v_j 1_{E_1}]$; such events exist by repeated applications of convex-rangedness. Denote $Y_j := v_j 1_{E_1} + \sum_{2\leq i\leq n} y_i 1_{E_i}$. Note that by construction, $A_j + B \in [Y_j]$ and $A + B \in [Y]$. By the event-continuity of ρ, $\lim_{j\to\infty} \rho(A_j + B) = \rho(A + B)$, and therefore

$$\lim_{j\to\infty} \widehat{\rho}(Y_j) = \lim_{j\to\infty} \rho(A_j + B) = \rho(A + B) = \widehat{\rho}(Y).$$

Likewise, take a decreasing sequence $\{v'_j\}$ in $\mathbf{D}$ converging to y_1, and denote $Y'_j := v'_j 1_{E_1} + \sum_{2\leq i\leq n} y_i 1_{E_i}$.. The same argument establishes that

$$\lim_{j\to\infty} \widehat{\rho}\left(Y'_j\right) = \widehat{\rho}(Y).$$

By the induction assumption, for all j,

$$\widehat{\rho}(Y_j) = I(Y_j) \text{ and } \widehat{\rho}\left(Y'_j\right) = I\left(Y'_j\right).$$

Hence, by the monotonicity of I,

$$\widehat{\rho}(Y) = \lim_{j\to\infty} \widehat{\rho}(Y_j) = \lim_{j\to\infty} I(Y_j) \leq I(Y) \leq \lim_{j\to\infty} I\left(Y'_j\right) = \lim_{j\to\infty} \widehat{\rho}\left(Y'_j\right) = \widehat{\rho}(Y),$$

which yields

$$\widehat{\rho}(Y) = I(Y)$$

as desired. $\square$

Proof of Theorem 2.

Step 1. constant-linearity of $\widehat{\rho}$ implies P4.

Take any $A, B \in \Sigma$ such that $\rho(A) \geq \rho(B)$, and any $x, y \in X$ with $u(y) < u(x)$. By b.u.s., we need to show that $\widehat{\rho}(u \circ [x_2, A; y, A^c]) \geq \widehat{\rho}(u \circ [x_2, B; y, B^c])$ .Indeed, this follows easily from the equalities $u \circ [x, A; y, A^c] = u(x)1_A + u(x)1_{A^c} = (u(x) - u(y))\, 1_A + u(y)\mathbf{1}$, whence by constant-linearity

$$\widehat{\rho}(u \circ [x, A; y, A^c]) = (u(x) - u(y))\, \rho(A) + u(y),$$

and similarly

$$\widehat{\rho}(u \circ [x, B; y, B^c]) = (u(x) - u(y))\, \rho(B) + u(y),$$

from which the desired conclusion follows immediately.

Step 2. P4 implies Union and Splitting Invariance.

Consider any $A \in \Sigma$, $\alpha, \beta \in [0, 1]$ such that $\alpha + \beta \leq 1$, and $A' \in \Lambda_A$ as well as $B_1 \in \Lambda_A$ and $B_2 \in \Lambda_{A^c}$ (both disjoint from A') such that $\overline{\pi}(A'/A) = \alpha$ and $\overline{\pi}(B_1/A) = \overline{\pi}(B_2/A^c) = \beta$. It clearly suffices to show that $\rho(A' + B) = \alpha\rho(A) + \beta$. Pick consequences y, x such that $u(y) = \beta$ and $u(x) = \alpha + \beta$. By utility sophistication and the conditional linearity property of $\widehat{\rho}$ (Lemma 1), $[x, A; y, A^c] \sim [x^+, A' + B_1; x^-, A \backslash (A' + B_1); x^+, ; x^-, A^c \backslash B_2]$, with the act on the right-hand side simplifying to $[x^+, A' + B; x^-, (A' + B)^c]$.

Moreover, taking any $T \in \Lambda$ with $\overline{\pi}(T) = \rho(A)$, by P4, $[x, A; y, A^c] \sim [x, T; y, T^c]$, and thus $[x^+, A' + B; x^-, (A' + B)^c] \sim [x, T; y, T^c]$. One computes $\widehat{\rho}(u \circ [x, T; y, T^c]) = E_{\overline{\pi}}(u \circ [x, T; y, T^c]) = (\alpha + \beta)\, \overline{\pi}(T) + \beta\pi^0(T^c) = \alpha\rho(A) + \beta$, whence

$$\rho(A' + B) = \widehat{\rho}(1_{A'+B}) = \widehat{\rho}(u \circ [x, T; y, T^c]) = \alpha\rho(A) + \beta.$$

Step 3a) Union Invariance implies constant-additivity.

Take any $Y = \sum_{i \in I} y_i 1_{E_i}$. Since $Y \leq (1 - c)\mathbf{1}$, there exist $A \in [Y]$ and $S, T \in \Lambda$ 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 $A_i \in \Lambda_{E_i}$ and $\overline{\pi}(A_i/E_i) = y_i$, $S = \sum_{i \in I} S_i$ with $S_i \in \Lambda_{E_i}$ and $\overline{\pi}(S_i/E_i) = \rho(A)$, and $T = \sum_{i \in I} T_i$ with $T_i \in \Lambda_{E_i}$ and $\overline{\pi}(S_i/E_i) = c$ 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 convex-rangedness of Π. Clearly, $A + T \in [Y + c\mathbf{1}]$. Since

$A \sim S$ by assumption, $A+T \sim S+T$ by Union Invariance; which is tantamount to $\rho(A+T) = \rho(S+T) = \rho(S) + \rho(T) = \rho(A) + c$. Hence

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

Step 3b) Splitting Invariance implies positive homogeneity

Take $Y \in \mathcal{Z}$ and rational $c = \frac{m}{n} \leq 1$, where m and n are natural numbers. Take $A \in [Y]$ and $T \in \Lambda$ such that $\overline{\pi}(T) = \widehat{\rho}(Y)$. By convex-rangedness, there exist partitions of A and T can be split into n equally likely subevents $\{A_1, ..., A_n\}$ and $\{T_1, ..., T_n\}$; by an argument paralleling that in i), the A_i can be chosen to belong to $[\frac{1}{n}Y]$, whence $\sum_{i \leq m} A_i \in [\frac{m}{n}Y]$. Since by construction $A \sim T$, by Splitting Invariance $A_1 \sim T_1$, and therefore by Splitting Invariance again $\sum_{i \leq m} A_i \sim \sum_{i \leq m} T_i$. It follows that

$$\widehat{\rho}\left(\frac{m}{n}Y\right) = \rho\left(\sum_{i \leq m} A_i\right) = \overline{\pi}\left(\sum_{i \leq m} T_i\right) = \frac{m}{n}\overline{\pi}(T) = \frac{m}{n}\widehat{\rho}(Y),$$

which establishes positive homogeneity for rational α. This implies positive homogeneity for arbitrary α, since by monotonicity of $\widehat{\rho}$,

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

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

Proof of Fact 2.

That ii) implies i) is straightforward.

The converse follows easily from showing that if $\succsim$ is strongly compatible with a complete $\trianglerighteq$ represented by a convex-ranged prior π, then $f \succsim g$ whenever $u \circ f$ stochastically dominates $u \circ g$. But viewing $\trianglerighteq$ as the revealed likelihood relation asssociated with $\succsim$, Strongly Compatibility is just Machina-Schmeidler's (1992) Strong Comparative Probability axiom, and the present claim is the key step in their proof. □

Proof of Proposition 3.

In view of Proposition 1, it is w.l.o.g. to assume that $X = [0,1]$ and $u(x) = x$ for all $x \in X$. Thus acts $f \in \mathcal{F}$ can be identified with random variables $Z \in \mathcal{Z}$.

1. Tradeoff Consistency of $\succsim$ implies Monotonicity and Lottery Independence of $\succsim_{AA}$

Lottery Independendence follows immediately from EU maximization over unambiguous acts implied by Theorem 1. To verify Monotonicity, take any act $F \in \mathcal{F}^{AA}$, lotteries $p, q \in \mathcal{L}$ and event

$S \in \Sigma_1$ such that $p \succsim_{AA} q$. By EU maximization over unambiguous acts, $p \sim^{AA} Ep$ and $q \sim^{AA} Eq$, hence $Ep \geq Eq$. Let Y and Z be Savage acts such that $F(Y) = [p$ on $S; F(\omega)$ elsewhere], $F(Z) = [q$ on $S; F(\omega)$ elsewhere], and $Y_{-S\times\Omega_2} = Z_{-S\times\Omega_2}$. By the Conditional Linearity property of the intrinsic integral (Lemma 1) and Theorem 1, $Y \sim [Ep$ on $S \times \Omega_2, Y_{-S\times\Omega_2}]$ and $Z \sim [Eq$ on $S \times \Omega_2, Z_{-S\times\Omega_2}]$. Since $Ep \geq Eq$, by Event-Monotonicity and transitivity therefore $Y \succsim Z$, which implies that [p on $S; F(\omega)$ elsewhere] $\succsim_{AA}$ [q on $S; F(\omega)$ elsewhere], as needs to be shown.

2. Monotonicity and Lottery Independence of $\succsim_{AA}$ imply Tradeoff Consistency of $\succsim$.

In view of Theorem 1, we need to show that $I = \widehat{\rho}$. That is, taking any $Z \in \mathcal{Z}$, we need to show that there exists $A \in [Z]$ such that $I(Z) = \rho(A)$, i.e. that $Z \sim 1_A$. To do this, write $Z = \sum_{i\leq n, j\leq n_i} z_{i,j} 1_{S_i\times T_{i,j}}$. By convex-rangedness, there exist for each $i \leq n, j \leq n_i$ events $T'_{i,j} \in \Lambda$ such that $T'_{i,j} \subseteq T_{i,j}$ and $\eta\left(T'_{i,j}\right) = z_{i,j}\eta\left(T_{i,j}\right)$. Let $A := \sum_{i\leq n, j\leq n_i} T'_{i,j}$. By construction, for all i, j, $A \cap S_i \times T_{i,j} = T'_{i,j}$ and $\pi\left(T'_{i,j}/S_i \times T_{i,j}\right) = z_{i,j}$, whence $A \in [Z]$. Write $F(Z) = [p_i$ on $S_i]$ and $F(1_A) = [q_i$ on $S_i]$. By the normalization of u,

$$z_{i,j} \sim^{AA} q_{ij} \text{ for all } i, j,$$

where $q_{ij} \in \mathcal{L}$ is defined by $q_{ij}^{x^+} = z_{i,j}$ and $q_{ij}^{x^-} = 1 - z_{i,j}$. Hence, noting that $q_i^{x^+} = Ep$, by Lottery Independence,

$$p_i \sim^{AA} q_i \text{ for all } i.$$

By Monotonicity, this implies

$$F(Z) \sim^{AA} F(1_A)$$

which yields $Z \sim 1_A$ as desired by the construction of $\succsim_{AA}$.

3. Tradeoff Consistency and P4 of $\succsim$ imply Certainty Independence of $\succsim_{AA}$.

Take any AA-acts $F = [p_i$ on $S_i], G = [q_i$ on $S_i]$, constant act $H = [q$ on $\Omega_1]$ and $\alpha \in (0, 1]$.By part 1), $F \sim^{AA} [Ep_i$ on $S_i]$, $G \sim^{AA} [Eq_i$ on $S_i]$, $\alpha F + (1-\alpha)H \sim^{AA} [\alpha Ep_i + (1-\alpha)Eq$ on $S_i]$, and $\alpha G + (1-\alpha)H \sim^{AA} [\alpha Eq_i + (1-\alpha)Eq$ on $S_i]$. Let $Y := \sum_i Ep_i 1_{S_i\times\Omega_2}$ and $Z := \sum_i Eq_i 1_{S_i\times\Omega_2}$. Clearly $F = F(Y)$, $G = F(Z)$, $\alpha F + (1-\alpha)H = F(\alpha Y + (1-\alpha)Eq)$, and $\alpha G + (1-\alpha)H = F(\alpha Z + (1-\alpha)Eq)$. By Theorem 2, $\widehat{\rho}$ is constant-linear. Hence by Theorem 1,

$$Y \succsim Z \text{ if and only if } \alpha Y + (1-\alpha)Eq \succsim \alpha Z + (1-\alpha)Eq.$$

In view of the equivalences established above, this translates back into the desired conclusion

$$F \succsim_{AA} G \text{ if and only if } \alpha F + (1-\alpha)H \succsim_{AA} \alpha G + (1-\alpha)H.$$

4. Monotonicity and Certainty Independence of $\succsim_{AA}$ imply P4 of $\succsim$.

In view of part 2) and Theorems 1 and 2, it suffices to establish the constant-linearity of the intrinsic integral $\widehat{\rho}$. Take any $Z \in \mathcal{Z}$, and $\alpha, c \in [0,1]$ such that $\alpha + c \leq 1$; we need to show that $\widehat{\rho}(\alpha Z + c\mathbf{1}) = \alpha\widehat{\rho}(Z + c)$ in any such case.

By definition,

$$Z \sim \widehat{\rho}(Z)\,\mathbf{1},$$

hence

$$F(Z) \sim^{AA} F(\widehat{\rho}(Z)\,\mathbf{1}) = [\widehat{\rho}(Z) \text{ on } \Omega_1].$$

By Certainty Independence, letting $H := F(\frac{c}{1-\alpha}\mathbf{1}) = [\frac{c}{1-\alpha} \text{ on } \Omega_1]$,

$$F(\alpha Z + c\mathbf{1}) = \alpha F + (1-\alpha)H \sim^{AA} \alpha[\widehat{\rho}(Z) \text{ on } \Omega_1] + (1-\alpha)H = F\left(\left(\alpha\widehat{\rho}(Z) + (1-\alpha)\frac{c}{1-\alpha}\right)\mathbf{1}\right),$$

whence from the definition of $\succsim_{AA}$,

$$\alpha Z + c\mathbf{1} \sim (\alpha\widehat{\rho}(Z) + c)\,\mathbf{1},$$

and therefore

$$\widehat{\rho}(\alpha Z + c\mathbf{1}) = \alpha\widehat{\rho}(Z) + c$$

by the normalization of $\widehat{\rho}$. $\square$

Proof of Proposition 4.

Suppose that a CEU preference ordering $\succsim$ is utility-sophisticated relative to the convex-ranged context $\trianglerighteq$. Since by EU maximization on unambiguous acts, $\nu = \rho$, we need to show that ρ is additive. Thus, take any disjoint events $A, B \in \Sigma$ as well as events $A' \subseteq A$ such that $A' \equiv A \backslash A'$, and $B' \subseteq B$ such that $B' \equiv B \backslash B'$. Specify consequences in utiles, and, for $z \in [0,1]$ let

$$f_z := [1 \text{ on } B', z \text{ on } A, 0 \text{ elsewhere}],$$

and

$$g_z := [\frac{1}{2} \text{ on } B, z \text{ on } A, 0 \text{ elsewhere}].$$

By construction, for all $\pi \in \Pi$, $E_\pi f_z = E_\pi g_z$, hence for all $z \in [0,1]$,

$$f_z \sim g_z$$

by utility-sophistication.

Now

$$\int u \circ f_z d\nu = \rho(B') + z[\rho(A+B') - \rho(B')],$$

while, for $z \geq \frac{1}{2}$

$$\int u \circ g_z d\nu = z\rho(A) + \frac{1}{2}[\rho(A+B) - \rho(A)],$$

while for $z \leq \frac{1}{2}$

$$\int u \circ g_z d\nu = \frac{1}{2}\rho(B) + z[\rho(A+B) - \rho(B)].$$

Thus, $\int u \circ f_z d\nu = \int u \circ g_z d\nu$ for $z \in \{0, \frac{1}{2}, 1\}$ only if

$$\rho(A) = \rho(A+B') - \rho(B') = \rho(A+B) - \rho(B),$$

i.e. only if $\rho(A) + \rho(B) = \rho(A+B)$, as needed to be shown. $\square$.

Proof of Proposition 5.

Part 1). Take any $x, y, x', y' \in X$ such that $x \succsim x'$, acts $f, g \in \mathcal{F}$ and events A disjoint from B and A' disjoint from B' such that $A \equiv B \rhd \emptyset$, $A' \equiv B'$ and [x on $A; y$ on $B; f(\omega)$ elsewhere] $\succsim$ [x' on $A; y'$ on $B; f(\omega)$ elsewhere]. By the convex-rangedness of $\trianglerighteq^0$, there exist events A'' disjoint from B'' and A''' disjoint from B''' such that $A'' + B'' = A + B$, $A''' + B''' = A' + B'$, $A'' \equiv B'' \rhd \emptyset$ and $A' \equiv B'$. Clearly by coherence and the fact that $\trianglerighteq \supseteq \trianglerighteq^0$, $B \backslash B'' = A'' \backslash A \equiv A \backslash A'' = B'' \backslash B$. Therefore by Strong Compatibility,

$$[x \text{ on } A; y \text{ on } B; f(\omega) \text{ elsewhere}] \sim [x \text{ on } A''; y \text{ on } B''; f(\omega) \text{ elsewhere}],$$

and

$$[x' \text{ on } A; y' \text{ on } B; f(\omega) \text{ elsewhere}] \sim [x' \text{ on } A''; y' \text{ on } B''; f(\omega) \text{ elsewhere}],$$

whence by transitivity also

$$[x \text{ on } A''; y \text{ on } B''; f(\omega) \text{ elsewhere}] \succsim [x' \text{ on } A''; y' \text{ on } B''; f(\omega) \text{ elsewhere}].$$

By Tradeoff Consistency with respect to $\trianglerighteq^0$, therefore

$$[x \text{ on } A'''; y \text{ on } B'''; f(\omega) \text{ elsewhere}] \succsim [x' \text{ on } A'''; y' \text{ on } B'''; f(\omega) \text{ elsewhere}].$$

By the same token as above,

$$[x \text{ on } A'''; y \text{ on } B'''; f(\omega) \text{ elsewhere}] \sim [x \text{ on } A'; y \text{ on } B'; f(\omega) \text{ elsewhere}],$$

and

$$[x' \text{ on } A'''; y' \text{ on } B'''; f(\omega) \text{ elsewhere}] \sim [x' \text{ on } A'; y' \text{ on } B'; f(\omega) \text{ elsewhere}],$$

whence by transitivity also

$$[x \text{ on } A'; y \text{ on } B'; f(\omega) \text{ elsewhere}] \succsim [x' \text{ on } A'; y' \text{ on } B'; f(\omega) \text{ elsewhere}]$$

as desired.

Part 2). Suppose $\succsim$ is utility sophisticated, Archimedean, bounded, solvable with respect to the convex-ranged $\trianglerighteq^0 \triangleright$. Clearly, $\succsim$ is also Archimedean with respect to $\trianglerighteq$, since $\triangleright$ contains $\triangleright^0$; likewise, $\succsim$ is solvable with respect to $\trianglerighteq$. By Theorem 1 ($2 \Longrightarrow 1$), preferences are tradeoff consistent with respect to $\trianglerighteq^0$. Since they are strongly compatible with $\trianglerighteq$ by assumption, they are also tradeoff consistent with respect to $\trianglerighteq$ by part 1) . By Theorem 1 again ($1 \Longrightarrow 2$), preferences are utility sophisticated with respect to $\trianglerighteq$. □

Proof of Fact 3.

W.l.o.g. assume that $\trianglerighteq^1$ is convex-ranged, with $\Pi^1 = co\{\mu_i\}_{i=1,\ldots,n}$. Take any $E \in \Sigma$ such that E is non-null with respect to both $\trianglerighteq^1$ and $\trianglerighteq^2$; otherwise there is nothing to prove.

Lemma 3 *Λ^2_E contains a σ-algebra in E Σ_E such that $\overline{\pi}^2_{\Sigma_E}$ defined as the restriction of $\overline{\pi}^2_{/E}$ to Σ_E is σ-additive and non-atomic.*

To verify the Lemma, let $\widetilde{\Pi}$ denote a dense set of σ-additive priors in Π^2, and $\mathcal{A}$ denote any algebra in E of subsets of Λ_E such that $\{\overline{\pi}^2(A/E) \mid A \in \mathcal{A}\}$ is dense in $[0,1]$. With Σ_E denoting the σ-algebra (in E) generated by $\mathcal{A}$, it follows from the uniqueness part of Caratheodory's Extension Theorem that, for all $A \in \Sigma_E$ and all $\pi, \pi' \in \widetilde{\Pi}$, $\pi(A/E) = \pi'(A/E)$. By the denseness of $\widetilde{\Pi}$ in Π^2 in the product topology, this conclusion extends to all $\pi, \pi' \in \Pi^2$, establishing that $\Sigma_E \subseteq \Lambda^2_E$. The σ-additivity of $\overline{\pi}^2_{\Sigma_E}$ is immediate, its non-atomicity follows from the dense-rangedness of $\overline{\pi}^2(./E)$ on $\mathcal{A}$ in view of Caratheodory's extension procedure (see, for example, Aliprantis-Border (1999, p. 343).

W.l.o.g. E is non-null for $\{\mu_i\}_{i=1,\ldots,m}$ with $m \geq 1$. For $i \leq m$, let $\mu_{i/E} = \frac{1}{\mu_i(E)}\mu_{i|\Sigma_E}$. restricted to Σ_E

Lemma 4 *For all $i \leq m$, $\mu_{i/E}$ is non-atomic.*

To see this by contradiction, assume that for some $i \leq m$, $\mu_{i/E}$ has an atom A in Σ_E. Since $\overline{\pi}^2_{\Sigma_E}$ is non-atomic by Lemma 3, there exists a decreasing sequence of events A_i in Σ_E such that $\mu_{i/E}(A_i) = \mu_{i/E}(A)$ and $\overline{\pi}^2_{\Sigma_E}(A_i) \downarrow 0$. By the σ-additivity of $\mu_{i/E}$, $\mu_{i/E}(\cap_i A_i) = 0$, while $\overline{\pi}^2_{\Sigma_E}(\cap_i A_i)$; thus $\cap_i A_i$ is $\trianglerighteq^2$-null but $\trianglerighteq^1$-non-null, in contradiction to the assumption of identical null sets.

By Lemmas 3 and 4, the measures $\mu_{i/E}$ (for $i \leq m$) and $\overline{\pi}^2_{\Sigma_E}$ are σ-additive and non-atomic. By Lyapunov's theorem, for all $\alpha \in [0,1]$, there exists $A \in \Sigma_E$ such that, for all $i \leq m$, $\mu_{i/E}(A) = \overline{\pi}^2_{\Sigma_E}(A) = \alpha$, from which the convex-rangedness of $co(\Pi^1 \cup \Pi^2) = \Pi_{(\trianglerighteq^1 \cap \trianglerighteq^2)}$ and thus of $\trianglerighteq^1 \cap \trianglerighteq^2$ follows easily. $\square$

REFERENCES

[1] Aliprantis, C. and K. Border (1999,): *Infinite Dimensional Analyis,* 2nd ed. Berlin: Springer.

[2] 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.

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

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

[5] Broome, J. (1991), *Weighing Goods,* Basil Blackwell, Oxford.

[6] Casadesus, R. , P. Klibanoff and E. Ozdenoren (2000): "Maxmin Expected Utility over Savage Acts with a Set of Priors", *Journal of Economic Theory* 92, 35-65.

[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. and M. Le Breton (1993): "Dynamically Consistent Beliefs must be Bayesian", *Journal of Economic Theory 63, 1-22.*

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

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

[12] Gilboa, I. (1987): "Expected Utility with Purely Subjective Nonadditive Probabilities," *Journal of Mathematical Economics* 16, 65-88.

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

[14] Ghirardato, P. (1997): "On Independence for Non-Additive Measures, with a Fubini Theorem," *Journal of Economic Theory* 73, 261-291.

[15] Ghirardato, P. and M. Marinacci (2001): "Risk, Ambiguity, and the Separation of Utility and Beliefs", *Mathematics of Operations Research* 26, 864-890 .

[16] Ghirardato, P. and M. Marinacci (2002a): "Ambiguity Made Precise: A Comparative Foundation" *Journal of Economic Theory* 102, 251–289.

[17] Ghirardato, P., F. Maccheroni and M. Marinacci (2002b): "Ambiguity from the Differential Viewpoint," ICER working paper 17/02.

[18] Ghirardato, P., F. Maccheroni, M. Marinacci and M. Siniscalchi (2003): "A Subjective Spin on Roulette Wheels", forthcoming *Econometrica*.

[19] Jaffray, J.-Y. (1989): "Linear Utility Theory for Belief Functions," *Operations Research Letters*, 9, pp. 107-112.

[20] Karni, E. (1996): "Probabilities and Beliefs", *Journal of Risk and Uncertainty* 13, 249-262.

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

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

[23] Klibanoff, P., Marinacci, M. and S. Mukerji (2002): "A Smooth Model of Decision Making Under Ambiguity", mimeo.

[24] Kopylov, I. (2002): "α-Maxmin expected utility" mimeo, University of Rochester.

[25] Luce, D. and H. Raiffa (1957): *Games and Decisions*, Dover.

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

[27] Machina, M. (2001): "Almost Objective Uncertainty", working paper UC San Diego.

[28] Machina, M. (2002): "Robustifying the Classical Model of Risk-Preferences and Beliefs", working paper UC San Diego.

[29] Marinacci, M., Maccheroni, F., Chateauneuf, A. and J.-M. Tallon (2003): "Monotone Continuous Multiple Priors", mimeo.

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

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

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

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

[34] Nehring, K. (2001): "Ambiguity in the Context of Probabilistic Beliefs", mimeo UC Davis.

[35] Nehring, K. (2003): "Imprecise Probabilistic Beliefs", Institute for Advanced Study WP #34.

[36] Nehring, K. (2004): "Ellsbergian Uncertainty Aversion", in preparation (expands material from Nehring (2001))

[37] 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.

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

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

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

[41] Siniscalchi, M. (2003): "A Behavioral Characterization of Plausible Priors", mimeo Northwestern.

[42] Tversky, A. and D. Kahneman (1992): "Advances in Prospect Theory: Cumulative Representations of Uncertainty," *Journal of Risk and Uncertainty* 5, 297-323.

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

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