DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES
CALIFORNIA INSTITUTE OF TECHNOLOGY
PASADENA, CALIFORNIA 91125

AMBIGUITY MADE PRECISE: A COMPARATIVE FOUNDATION

Paolo Ghirardato
California Institute of Technology

Massimo Marinacci
Università di Torino

SOCIAL SCIENCE WORKING PAPER 1026

October 1997
Revised: December 2000

Ambiguity Made Precise: A Comparative Foundation

Paolo Ghirardato Massimo Marinacci

Abstract

The theory of subjective expected utility (SEU) has been recently extended to allow ambiguity to matter for choice. We propose a notion of absolute ambiguity aversion by building on a notion of comparative ambiguity aversion. We characterize it for a preference model which encompasses some of the most popular models in the literature. We next build on these ideas to provide a definition of unambiguous act and event, and show the characterization of the latter. As an illustration, we consider the classical Ellsberg 3-color urn problem and find that the notions developed in the paper provide intuitive answers.

Journal of Economic Literature Classification Number: D81

JEL classification numbers: D81

Key words: Ambiguity Aversion, Cardinal Risk Aversion, Choquet Expected Utility, Maxmin Expected Utility, Biseparable Preferences

Ambiguity Made Precise: A Comparative Foundation*

Paolo Ghirardato Massimo Marinacci

Introduction

In this paper we propose and characterize a formal definition of ambiguity aversion for a class of preference models which encompasses the most popular models developed to allow ambiguity attitude in decision making. Using this notion, we define and characterize ambiguity of events for ambiguity averse or loving preferences. Our analysis is based on a fully 'subjective' framework with no extraneous devices (like a roulette wheel, or a rich set of exogenously 'unambiguous' events). This yields a definition that can be fruitfully used with any preference in the mentioned class, though it imposes a limitation in the definition's ability of distinguishing 'real' ambiguity aversion from other behavioral traits that have been observed experimentally.

The subjective expected utility (SEU) theory of decision making under uncertainty of Savage [25] is firmly established as the choice-theoretic underpinning of modern economic theory. However, such success has well known costs: SEU's simple and powerful representation is often violated by actual behavior, and it imposes unwanted restrictions. In particular, Ellsberg's [7] famous thought experiment (see Section 5) convincingly shows that SEU cannot take into account the possibility that the information a decision maker (DM) has about some relevant uncertain event is vague or imprecise, and that such 'ambiguity' affects her behavior. Ellsberg observed that ambiguity affected his 'nonexperimental' subjects in a consistent fashion: Most of them preferred to bet on unambiguous rather than ambiguous events. Furthermore, he found that even when shown the inconsistency of their behavior with SEU, the subjects stood their ground "because it seems to them the sensible way to behave." This attitude has later been named *ambiguity aversion*,

* An earlier version of this paper was circulated with the title "Ambiguity Made Precise: A Comparative Foundation and Some Implications". We thank Kim Border, Eddie Dekel, Itzhak Gilboa, Tony Kwasnica, Antonio Rangel, David Schmeidler, audiences at Caltech, Johns Hopkins, Northwestern, NYU, Rochester, UC-Irvine, Université Paris I, the TARK VII-Summer Micro Conference (Northwestern, July 1998), the 1999 RUD Workshop, and especially Simon Grant, Peter Klibanoff, Biung-Ghi Ju, Peter Wakker, and an anonymous referee for helpful comments and discussion. Our greatest debt of gratitude is however to Larry Epstein, who sparked our interest on this subject with his paper [8], and stimulated it with many discussions. Marinacci gratefully acknowledges the financial support of MURST.

and has received ample experimental confirmation.[1] Savage was well aware of this limit of SEU, for he wrote that

> [...] There seem to be some probability relations about which we feel relatively 'sure' as compared with others. [...] The notion of 'sure' and 'unsure' introduced here is vague, and my complaint is precisely that neither the theory of personal probability, as it is developed in this book, nor any other device known to me renders the notion less vague. [25, pp. 57–58 of the 1972 edition]

In the wake of Ellsberg's contribution, extensions of SEU have been developed allowing ambiguity, and the DM's attitude towards it, to play a role in her choices. Two methods for extending SEU have established themselves as the standards of this literature. The first, originally proposed in Schmeidler [26], is to allow the DM's beliefs on the state space to be represented by non-additive probabilities, called *capacities*, and her preferences by *Choquet integrals* (which are just standard integrals when integrated with respect to additive probabilities). For this reason, this generalization is called the theory of *Choquet expected utility* (CEU) maximization. The second, axiomatized by Gilboa and Schmeidler [15], allows the DM's beliefs to be represented by multiple probabilities, and represents her preferences by the 'maximin' on the set of the expected utilities. This generalization is thus called the *maxmin expected utility* (MEU) theory. Here we use the general class of preferences with ambiguity attitudes developed in our [13]. These orderings, that we call *biseparable preferences*, are all those such that the ranking of consequences can be represented by a state-independent cardinal utility u, and the ranking of *bets on events* by u and a unique numerical function (a capacity) ρ.[2] The latter represents the DM's *willingness to bet*; i.e., $\rho(A)$ is roughly the number of euros she is willing to exchange for a bet that pays 1 euro if event A obtains and 0 euros otherwise. The only restriction imposed on the ranking of non-binary acts is a mild dominance condition. CEU and MEU are special cases of biseparable preferences, where ρ is respectively the DM's non-additive belief and the lower envelope of her multiple probabilities.

An important reason for the lasting success of SEU theory is the elegant theory of the measurement of risk aversion developed from the seminal contributions of de Finetti [6], Arrow [2] and Pratt [24]. Unlike risk aversion, ambiguity aversion is yet without a fully general formalization, one that does not require extraneous devices and applies to most if not all the existing models of ambiguity averse behavior. This paper attempts to fill this gap: We propose a definition of ambiguity aversion and show its formal characterization in the general decision-theoretic framework of Savage, whose only restriction is a richness condition on the set of consequences. Our definition is *behavioral*; that is, it only requires observation of the DM's preferences on acts in this fully subjective setting. However, the

[1] Other widespread names are 'uncertainty aversion' and 'aversion to Knightian uncertainty'. We like to use 'uncertainty' in its common meaning of *any* situation in which the consequences of the DM's possible actions are not known at the time of choice.

[2] A bet 'on' an event is any binary act in which a better payoff ('win') is received when the event obtains.

definition works as well (indeed better, see Proposition 11) in the *Anscombe-Aumann framework*, a special case of Savage's framework which presumes the existence of an auxiliary device with 'known' probabilities.

Decision models with ambiguity averse preferences the objects of increasing attention by economists and political scientists interested in explaining phenomena at odds with SEU. For example, they have been used to explain the existence of incomplete contracts (Mukerji [22]), the existence of substantial volatility in stock markets (Epstein and Wang [9], Hansen, Sargent and Tallarini [16]), or selective abstention in political elections (Ghirardato and Katz [12]). We hope that the characterization provided here will turn out to be useful for the 'applications' of models of ambiguity aversion, as that of risk aversion was for the applications of SEU. More concretely, we hope that it will help to understand the predictive differences of risk and ambiguity attitudes.

To understand our definition, it is helpful to go back to the characterization of risk aversion in the SEU model. The following approach to defining risk aversion was inspired by Yaari [30]. Given a state space S, let $\mathcal{F}$ denote a collection of 'acts', maps from S into $\mathbb{R}$ (e.g., monetary payoffs). Define a *comparative* notion of risk aversion for SEU preferences as follows: Say that $\succcurlyeq_2$ is *more risk averse than* $\succcurlyeq_1$ if they have identical beliefs and the following implications hold for every 'riskless' (i.e., constant) act x and every 'risky' act f:

$$x \succcurlyeq_1 f \quad \Rightarrow \quad x \succcurlyeq_2 f \tag{1}$$

$$x \succ_1 f \quad \Rightarrow \quad x \succ_2 f \tag{2}$$

(where $\succ$ is the asymmetric component of $\succcurlyeq$). Identity of beliefs is required to avoid possible confusions between differences in risk attitudes and in beliefs (cf. Yaari [30, p.317]). We can use this comparative ranking to obtain an *absolute* notion of risk aversion by calling some DMs — for instance expected value maximizers — *risk neutral*, and by then calling *risk averse* those DMs who are more risk averse than risk neutrals. As it is well known, this 'comparatively founded' notion has the usual characterization. Like the traditional 'direct' definition of risk aversion, it is fully behavioral in the sense defined above. However, its interpretation is based on two primitive assumptions. First, constant acts are *intuitively* riskless. Second, expected value maximization *intuitively* reflects risk neutral behavior, so that it can be used as our benchmark for measuring risk aversion.

In this paper, we follow the example of Epstein [8] in giving a comparative foundation to ambiguity attitude: We start from a 'more ambiguity averse than...' ranking and then establish a benchmark, thus obtaining an 'absolute' definition of ambiguity aversion. Analogously to Yaari's, our 'more ambiguity averse...' relation is based on the following intuitive consideration: *If a DM prefers an unambiguous (resp. ambiguous) act to an ambiguous (resp. unambiguous) one, a more (resp. less) ambiguity averse one will do the same.* This is natural, but it raises the obvious question of which acts should be used as the 'unambiguous' acts for this ranking. Depending on the decision problem the DM is facing and on her information, there might be different sets of 'obviously' unambiguous acts; i.e., acts that we are confident that any DM perceives as unambiguous. It seems intuitive to us that in any well-formulated problem, the constant acts will be in this set. Hence, we make our *first primitive assumption*: Constant acts are the only acts that are 'obviously' unambiguous in any problem, since other acts may not be perceived as unambiguous by some DM in some state of information. This assumption implies that a preference (not necessarily SEU) $\succcurlyeq_2$ is more ambiguity averse than $\succcurlyeq_1$ whenever Eqs. (1) and (2) hold. However, the following example casts some doubts as to the intuitive appeal of such definition:

Example 1 Consider an (Ellsberg) urn containing balls of two colors: Black and Red. Two DMs are facing this urn, and they have no information on its composition. The first DM has SEU preferences $\succcurlyeq_1$, with a utility function on the set of consequences $\mathbb{R}$ given by $u_1(x) = x$, and beliefs on the state space of ball extractions $S = \{B, R\}$ given by

$$\rho_1(B) = \frac{1}{2} \quad \text{and} \quad \rho_1(R) = \frac{1}{2}.$$

The second DM also has SEU preferences, and identical beliefs: Her preference $\succcurlyeq_2$ is represented by $u_2(x) = \sqrt{x}$ and $\rho_2 = \rho_1$. Both (1) and (2) hold, but it is quite clear that this is due to differences in the DMs' risk attitudes, and not in their ambiguity attitudes: They both apparently disregard the ambiguity in their information. $\triangle$

Given a biseparable preference, call *cardinal risk attitude* the psychological trait described by the utility function u — what explains any differences in the choices over bets of two biseparable preferences with the *same* willingness to bet ρ. The problem with the example is that the two DMs have different cardinal risk attitude. To avoid confusions of this sort, our comparative ambiguity ranking uses Eqs. (1) and (2) only on pairs which satisfy a behavioral condition, called *cardinal symmetry*, that implies that two DMs have identical u. As it only looks at each DM's preferences over bets on one event (which may be different across DMs), cardinal symmetry does not impose any restriction on the DMs' relative ambiguity attitudes.

Having thus constructed the comparative ambiguity ranking, we next choose a benchmark against which to measure ambiguity aversion. It seems generally agreed that SEU preferences are *intuitively* ambiguity neutral. We use SEU preferences as benchmarks because we posit — our *second primitive assumption* — that they are the only ones that are 'obviously' ambiguity neutral in any decision problem and in any situation. Thus, *ambiguity averse* is any preference relation $\succcurlyeq$ for which there is a SEU preference 'less ambiguity averse than' $\succcurlyeq$. Ambiguity love and (endogenous) neutrality are defined in the obvious way.

The main results in the paper present the characterization of these notions of ambiguity attitude for biseparable preferences. The characterization of ambiguity neutrality is simply stated: A preference is ambiguity neutral if and only if it has a SEU representation. That is, the *only* preferences which are endogenously ambiguity neutral are SEU. The general characterization of ambiguity aversion (resp. love) implies in particular that a preference is ambiguity averse (resp. loving) *only if* its willingness to bet ρ is pointwise dominated by (resp. pointwise dominates) a probability. In the CEU case, the converse is also true: A CEU preference is ambiguity averse *if and only if* its belief (which is equal to ρ) is dominated by a probability; i.e., it has a non-empty 'core'. On the other hand, all MEU preferences are ambiguity averse, as it is intuitive. As to comparative ambiguity aversion, we find that *if* $\succcurlyeq_2$ is more ambiguity averse than $\succcurlyeq_1$ *then* $\rho_1 \geq \rho_2$. That is, a less ambiguity averse DM will have uniformly higher willingness to bet. The latter condition is also sufficient for CEU preferences, whereas for MEU preferences containment of the sets of probabilities is necessary and sufficient for relative ambiguity.

We next briefly turn to the issue of defining ambiguity itself. A 'behavioral' notion of unambiguous act follows naturally from our earlier analysis: Say that an act is *unambiguous* if an ambiguity averse (or loving) DM evaluates it in an ambiguity neutral fashion. The unambiguous *events* are those that unambiguous acts depend upon. We obtain the following simple characterization of the set of unambiguous events for biseparable preferences: For an ambiguity averse (or loving) DM with willingness to bet ρ, event A is unambiguous if and only if $\rho(A) + \rho(A^c) = 1$. (A more extensive discussion of ambiguity is contained in the companion [14].)

Finally, as an application of the previous analysis, we consider the classical Ellsberg problem with a 3-color urn. We show that the theory delivers the intuitive answers, once the information provided to the DM is correctly incorporated.

It is important to underscore from the outset two important limitations of the notions of ambiguity attitude we propose. The first limitation is that while the comparative foundation makes our absolute notion notion 'behavioral', in the sense defined above, it also makes it *computationally* demanding. A more satisfactory definition would be one which is more 'direct': It can be verified by observing a smaller subset of the DM's preference relation. While we conjecture that it may be possible to construct such a definition — obtaining the same characterization as the one proposed here — we leave its development to future work.

Our comparative notion is more direct, thus less amenable to this criticism. However, it is in turn limited by the requirement of the identity of cardinal risk attitude. The absolute notion is not, as it conceptually builds on the comparison of the DM with an idealized version of herself, identical to her in all traits *but her ambiguity aversion.*

The second limitation stems from the fact that no extraneous devices are used in this paper. An advantage of this is that our notions apply to any decision problem under uncertainty, and our results to any biseparable preference. However, such wide scope carries costs: Our notion of ambiguity aversion comprises behavioral traits that may not be due to ambiguity — like *probabilistic risk aversion*, the tendency of discounting 'objective' probabilities that has been observed in many experiments on decision making under risk (including the celebrated 'Allais paradox'). Thus, one may consider it more appropriate to use a different name for what is measured here, like 'chance aversion' or 'extended ambiguity aversion'.

The reason for our choice of terminology is that we see a ranking of conceptual importance between ambiguity aversion/love and other departures from SEU maximization. As we argued above using Savage's words, the presence of ambiguity provides a *normatively* compelling reason for violating SEU. We do not feel that other documented reasons are similarly compelling. Moreover, we hold (see below and Subsection 6.3) that extraneous devices — say, a rich set of exogenously 'unambiguous' events — are required for ascertaining the reason of a given departure. Thus, when these devices are not available — say, because the set of 'unambiguous' events is not rich enough — we prefer to attribute a departure to the reasons we find normatively more compelling. However, the reader is warned, so that he/she may choose to give a different name to the phenomenon we formally describe.

The Related Literature

The problem of defining ambiguity and ambiguity aversion is discussed in a number of earlier papers. The closest to ours in spirit and generality is Epstein [8], the first paper to develop a notion of absolute ambiguity aversion from a comparative foundation.[3] As we discuss in more detail in Subsection 6.3, the comparative notion and benchmarks he uses are different from ours. Epstein's objective is to provide a more precise measurement of ambiguity attitude than the one we attempt here; in particular, to filter out probabilistic risk aversion. For this reason, he assumes that in the absence of ambiguity a DM's preferences are 'probabilistically sophisticated' in the sense of Machina and Schmeidler [20]. However, we argue that for its conclusions to conform with intuition, Epstein's approach requires an extraneous device: a rich set of acts which are exogenously established to be 'unambiguous', much larger than the set of the constants that we use. Thus, the higher accuracy of his approach limits its applicability *vis à vis* our cruder but less demanding approach.

The most widely known and accepted definition of absolute ambiguity aversion is that proposed by Schmeidler in his seminal CEU model [26]. Employing an Anscombe-Aumann framework, he defines ambiguity aversion as the preference for 'objective mixtures' of acts, and he shows that for CEU preferences this notion is characterized by the *convexity* of the capacity representing the DM's beliefs. While the intuition behind this definition is certainly compelling, Schmeidler's axiom captures *more* than our notion of ambiguity aversion. It gives rise to ambiguity averse behavior, but it entails additional structure that does not seem to be related to ambiguity aversion (see Example 25). Doubts about the relation of convexity to ambiguity aversion in the CEU case are also raised by Epstein [8], but he concludes that they are completely unrelated (see Section 5 for a discussion).

There are other interesting papers dealing with ambiguity and ambiguity aversion. In a finite setting, Kelsey and Nandeibam [18] propose a notion of comparative ambiguity for

[3] There are earlier papers that use a comparative approach for studying ambiguity attitude, but they do not use it as a basis for defining absolute notions. E.g., Tversky and Wakker [27].

the CEU and MEU models similar to ours and obtain a similar characterization, as well as an additional characterization in the CEU case. Unlike us, they do not consider absolute ambiguity attitude, and they do not discuss the issue of the distinction of cardinal risk and ambiguity attitude. Montesano and Giovannoni [21] notice a connection between absolute ambiguity aversion in the CEU model and nonemptiness of the core, but they base themselves purely on intuitive considerations on Ellsberg's example. Chateauneuf and Tallon [4] present an intuitive necessary and sufficient condition for non-emptiness of the core of CEU preferences in an Anscombe-Aumann framework. Zhang [31], Nehring [23], and Epstein and Zhang [10] propose different definitions of unambiguous event and act. Fishburn [11] characterizes axiomatically a primitive notion of ambiguity.

Organization

The structure of the paper is as follows. Section 1 provides the necessary definitions and set-up. Section 2 introduces the notions of ambiguity aversion. The cardinal symmetry condition is introduced in Subsection 2.1, and the comparative and absolute definitions in 2.2. Section 3 presents the characterization results. Section 4 contains the notions of unambiguous act and event, and the characterization of the latter. In Section 5, we go back to the Ellsberg urn and show the implications of our results for that example. Section 6 discusses the key aspects of our approach, in particular, the choices of the comparative ambiguity ranking and the benchmark for defining ambiguity neutrality; it thus provides a more detailed comparison with Epstein's [8] approach. The Appendices contain the proofs and some technical material.

1 Set-Up and Preliminaries

The general set-up of Savage [25] is the following. There is a set S of states of the world, an algebra Σ of subsets of S, and a set X of consequences. The choice set $\mathcal{F}$ is the set of all *finite-valued* acts $f : S \to X$ which are measurable w.r.t. Σ. With the customary abuse of notation, for $x \in X$ we define $x \in \mathcal{F}$ to be the constant act $x(s) = x$ for all $s \in S$, so that $X \subseteq \mathcal{F}$. Given $A \in \Sigma$, we denote by $x\,A\,y$ the binary act (bet) $f \in \mathcal{F}$ such that $f(s) = x$ for $s \in A$, and $f(s) = y$ for $s \notin A$.

Our definitions require that the DM's preferences be represented by a **weak order** $\succcurlyeq$ on $\mathcal{F}$: a complete and transitive binary relation $\succcurlyeq$, with asymmetric (resp. symmetric) component $\succ$ (resp. $\sim$). The weak order $\succcurlyeq$ is called **nontrivial** if there are $f, g \in \mathcal{F}$ such that $f \succ g$. We henceforth call **preference relation** any nontrivial weak order on $\mathcal{F}$.

A functional $V : \mathcal{F} \to \mathbb{R}$ is a **representation** of $\succcurlyeq$ if for every $f, g \in \mathcal{F}$, $f \succcurlyeq g$ if and only if $V(f) \geq V(g)$. A representation V is called: **monotonic** if $f(s) \succcurlyeq g(s)$ for every $s \in S$ implies $V(f) \geq V(g)$; **nontrivial** if $V(f) > V(g)$ for some $f, g \in \mathcal{F}$.

While the definitions apply to any preference relation, our results require a little more structure, provided by a general decision model introduced in Ghirardato and Marinacci [13]. To present it, we need the following notion of 'nontrivial' event: Given a preference relation $\succcurlyeq$, $A \in \Sigma$ is **essential for** $\succcurlyeq$ if for some $x, y \in X$, we have $x \succ x\,A\,y \succ y$.

Definition 2 *Let $\succcurlyeq$ be a binary relation. We say that a representation $V : \mathcal{F} \to \mathbb{R}$ of $\succcurlyeq$ is **canonical** if it is nontrivial monotonic and there exists a set-function $\rho : \Sigma \to [0,1]$ such that, letting $u(x) \equiv V(x)$ for all $x \in X$, for all consequences $x \succcurlyeq y$ and all events A,*

$$V(x\,A\,y) = u(x)\,\rho(A) + u(y)\,(1 - \rho(A)). \tag{3}$$

*A relation $\succcurlyeq$ is called a **biseparable preference** if it admits a canonical representation, and moreover such representation is unique up to a positive affine transformation when $\succcurlyeq$ has at least one essential event.*

Clearly, a biseparable preference is a preference relation. If V is a canonical representation of $\succcurlyeq$, then u is a cardinal state-independent representation of the DM's preferences over consequences, hence we call it his **canonical utility index**. Moreover, for all $x \succ y$ and all events $A, B \in \Sigma$ we have $x\,A\,y \succcurlyeq x\,B\,y$ if and only if $\rho(A) \geq \rho(B)$. Thus, ρ represents the DM's **willingness to bet** (likelihood relation) on events. Moreover, ρ is easily shown to be a **capacity** — a set-function normalized and monotonic w.r.t. set inclusion — so that V evaluates binary acts by taking the Choquet expectation of u with respect to ρ.[4] However, the DM's preferences over non-binary acts are not constrained to a specific functional form.

To understand the rationale of the clause relating to essential events, first observe that for any $\succcurlyeq$ with a canonical representation with willingness to bet ρ, an event A is essential if and only if $0 < \rho(A) < 1$. Thus, there are no essential events iff $\rho(A)$ is either 0 or 1 for every A; that is, the DM behaves as if he does not judge any bet to be uncertain, and his canonical utility index is ordinal. In such a case, the DM's cardinal risk attitude is then intuitively not defined: without an uncertain event there is no risk. On the other hand, it can be shown [13, Theorem 4] that cardinal risk attitude is characterized by a cardinal property of the canonical utility index, its concavity. Hence the additional requirement in Definition 2 guarantees that when there is some uncertain event cardinal risk aversion is well defined.

As the differences in two DM's cardinal risk attitude might play a role in the choices in Eqs. (1) and (2), it is useful to identify the situation in which these attitudes are defined: Say that preference relations $\succcurlyeq_1$ and $\succcurlyeq_2$ **have essential events** if there are events $A_1, A_2 \in \Sigma$ such that for each $i = 1, 2$, A_i is essential for $\succcurlyeq_i$.

To avoid repetitions, the following lists all the assumptions on the structure of the decision problem and on the DM's preferences *that are tacitly assumed in all results in the paper*:

[4] See Appendix A for the definition of capacities, Choquet integrals, and some of their properties.

Structural Assumption *X is a connected and separable topological space (e.g., a convex subset of $\mathbb{R}^n$ with the usual topology). Every biseparable preference on $\mathcal{F}$ has a continuous canonical utility function.*

A full axiomatic characterization of the biseparable preferences satisfying the Structural Assumption is provided in our [13].

1.1 Some Examples of Biseparable Preferences

As mentioned above, the biseparable preference model is very general. In fact, it contains most of the known preference models that obtain a separation between cardinal (state-independent) utility and willingness to bet. We now illustrate this claim by showing some examples of decision models which under mild additional restrictions (e.g., the Structural Assumption) belong to the biseparable class. (More examples and details are found in [13].)

(*i*) A binary relation $\succcurlyeq$ on $\mathcal{F}$ is a **CEU ordering** if there exist a cardinal utility index u on X and a capacity ν on (S, Σ) such that $\succcurlyeq$ can be represented by the functional $V : \mathcal{F} \to \mathbb{R}$ defined by the following equation:

$$V(f) = \int_S u(f(\cdot))\, d\nu, \tag{4}$$

where the integral is taken in the sense of Choquet (notice that it is finite because each act in $\mathcal{F}$ is finite-valued). The functional V is immediately seen to be a canonical representation of $\succcurlyeq$, and $\rho = \nu$ is its willingness to bet. An important subclass of CEU orderings are the **SEU orderings**, which correspond to the special case in which ν is a probability measure, i.e., a finitely additive capacity. See Wakker [28] for an axiomatization of CEU and SEU preferences (satisfying the Structural Assumption) in the Savage setting.

(*ii*) Let Δ denote the set of all the probability measures on (S, Σ). A binary relation $\succcurlyeq$ on $\mathcal{F}$ is a **MEU ordering** if there exist a cardinal utility index u and a unique non-empty, (weak*)-compact and convex set $C \subseteq \Delta$ such that $\succcurlyeq$ can be represented by the functional $V : \mathcal{F} \to \mathbb{R}$ defined by the following equation:

$$V(f) = \min_{P \in C} \int_S u(f(s))\, P(ds). \tag{5}$$

SEU also corresponds to the special case of MEU in which $C = \{P\}$ for some probability measure P. If we now let for any $A \in \Sigma$,

$$\underline{P}(A) = \min_{P \in C} P(A), \tag{6}$$

we see that $\underline{P}$ is an exact capacity. While in general $V(f)$ is *not* equal to the Choquet integral of $u(f)$ with respect to $\underline{P}$, this is the case for binary acts f. This shows that V is a canonical representation of $\succcurlyeq$, with willingness to bet $\rho = \underline{P}$. See Casadesus-Masanell, Klibanoff and Ozdenoren [3] for an axiomatization of MEU preferences (satisfying the Structural Assumption) in the Savage setting.

More generally, consider an **α-MEU** preference which assigns some weight to both the worst-case and best-case scenarios. Formally, there is a cardinal utility u, a set or probabilities C, and $\alpha \in [0,1]$, such that $\succcurlyeq$ is represented by

$$V(f) = \left[\alpha \min_{P\in C} \int_S u(f(s))\, P(ds) + (1-\alpha) \max_{P\in C} \int_S u(f(s))\, P(ds)\right].$$

This includes the case of a 'maximax' DM, who has $\alpha \equiv 0$. V is canonical, so that $\succcurlyeq$ is biseparable, with ρ given by $\rho(A) = \alpha \min_{P\in C} P(A) + (1-\alpha) \max_{P\in C} P(A)$, for $A \in \Sigma$.

(*iii*) Consider a binary relation $\succcurlyeq$ constructed as follows: There is a cardinal utility u, a probability P and a number $\beta \in [0,1]$ such that $\succcurlyeq$ is represented by

$$V(f) \equiv (1-\beta) \int_S u(f(s))\, P(ds) + \beta\, \varphi(u \circ f),$$

where

$$\varphi(u\circ f) \equiv \sup \left\{ \int_S u(g(s))\, P(ds) : g \in \mathcal{F} \text{ binary, } u(g(s)) \le u(f(s)) \text{ for all } s \in S \right\}.$$

$\succcurlyeq$ describes a DM who behaves as if he was maximizing SEU when choosing among binary acts, but not when comparing more complex acts. The higher the parameter β, the farther the preference relation is from SEU on non-binary acts. V is monotonic and it satisfies Eq. (3) with $\rho = P$, so that it is a canonical representation of $\succcurlyeq$.

1.2 The Anscombe-Aumann Case

The **Anscombe-Aumann framework** is a widely used special case of our framework in which the consequences have an objective feature: X is also a convex subset of a vector space. For instance, X is the set of all the lotteries on a set of prizes if the DM has access to an 'objective' independent randomizing device. In this framework, it is natural to consider the following variant of the biseparable preference model — where for every $f, g \in \mathcal{F}$ and $\alpha \in [0,1]$, $\alpha f + (1-\alpha)g$ denotes the act which pays $\alpha f(s) + (1-\alpha)g(s) \in X$ for every $s \in S$.

Definition 3 *A canonical representation V of a preference relation $\succcurlyeq$ is* ***constant linear (c-linear*** *for short) if $V(\alpha f + (1-\alpha)x) = \alpha V(f) + (1-\alpha)V(x)$ for all binary $f \in \mathcal{F}$, $x \in X$, and $\alpha \in [0,1]$. A relation is called a* ***c-linearly biseparable preference*** *if it admits a c-linear canonical representation.*

Again, an axiomatic characterization of this model is found in [13]. It generalizes the SEU model of Anscombe and Aumann [1] and many non-EU extensions that followed, like the CEU and MEU models of Schmeidler [26] and Gilboa and Schmeidler [15] respectively. In fact, a c-linearly biseparable preference behaves in a SEU fashion over the set X of the constant acts, but it is almost unconstrained over non-binary acts. (C-linearity guarantees the cardinality of V and hence u.)

All the results in this paper are immediately translated to this class of preferences, in particular to the CEU and MEU models in the Anscombe-Aumann framework mentioned above. Indeed, as we show in Proposition 11 below, in this case removing cardinal risk aversion is much easier than in the more general framework we use.

2 The Definitions

As anticipated in the Introduction, the point of departure of our search for an extended notion of ambiguity aversion is the following partial order on preference relations:

Definition 4 *Let $\succcurlyeq_1$ and $\succcurlyeq_2$ be two preference relations. We say that $\succcurlyeq_2$ is **more uncertainty averse than** $\succcurlyeq_1$ if: For all $x \in X$ and $f \in \mathcal{F}$, both*

$$x \succcurlyeq_1 f \Rightarrow x \succcurlyeq_2 f \tag{7}$$

and

$$x \succ_1 f \Rightarrow x \succ_2 f. \tag{8}$$

This order has the advantage of making the weakest prejudgment on which acts are 'intuitively' unambiguous: The constants. However, Example 1 illustrates that it does not discriminate between cardinal risk attitude and ambiguity attitude: DMs 1 and 2 are intuitively *both* ambiguity neutral, but 1 is more cardinal risk averse, and hence more uncertainty averse than 2. The problem is that constant acts are 'neutral' with respect to ambiguity *and* with respect to cardinal risk. Given that our objective is comparing ambiguity attitudes, we thus need to find ways to coarsen the ranking above, so as to identify which part is due to differences in cardinal risk attitude and which is due to differences in ambiguity attitude.

2.1 Filtering Cardinal Risk Attitude

While the 'factorization' just described can be achieved easily if we impose more structure on the decision framework (see, e.g., the discussion in Subsection 6.3), we present a method for separating cardinal risk and ambiguity attitude which is only based on preferences, does not employ extraneous devices, and obtains the result for all biseparable preferences. Moreover, this approach does not impose any restrictions on the two DMs' beliefs (and hence on their relative ambiguity attitude), a problem that all alternatives share. The key step is coarsening comparative uncertainty aversion by adding the following restriction on which pairs of preferences are to be compared (we write $\{x, y\} \succ z$ as a short-hand for $x \succ z$ and $y \succ z$, and similarly for $\prec$):

Definition 5 *Two preference relations $\succcurlyeq_1$ and $\succcurlyeq_2$ are **cardinally symmetric** if for any pair $(A_1, A_2) \in \Sigma \times \Sigma$ such that each A_i is essential for $\succcurlyeq_i$, $i = 1, 2$, and any $v_*, v^*, w_*, w^* \in X$ such that $v_* \prec_1 v^*$ and $w_* \prec_2 w^*$ we have:*

- *If there are $x, y \in X$ such that $v_* \succ_1 \{x, y\}$, $w_* \succ_2 \{x, y\}$, and*

$$v_* \, A_1 \, x \sim_1 v^* \, A_1 \, y \qquad \text{and} \qquad w_* \, A_2 \, x \sim_2 w^* \, A_2 \, y, \tag{9}$$

  then for every $x', y' \in X$ such that $v_* \succ_1 \{x', y'\}$, $w_* \succ_2 \{x', y'\}$ we have

$$v_* \, A_1 \, x' \sim_1 v^* \, A_1 \, y' \Longleftrightarrow w_* \, A_2 \, x' \sim_2 w^* \, A_2 \, y'. \tag{10}$$

- *Symmetrically, if there are $x, y \in X$ such that $v^* \prec_1 \{x, y\}$, $w^* \prec_2 \{x, y\}$, and*

$$x \, A_1 \, v^* \sim_1 y \, A_1 \, v_* \qquad \text{and} \qquad x \, A_2 \, w^* \sim_2 y \, A_2 \, w_*, \tag{11}$$

  then for every $x', y' \in X$ such that $v^* \prec_1 \{x', y'\}$, $w^* \prec_2 \{x', y'\}$ we have

$$x' \, A_1 \, v^* \sim_1 y' \, A_1 \, v_* \Longleftrightarrow x' \, A_2 \, w^* \sim_2 y' \, A_2 \, w_*. \tag{12}$$

This condition is inspired by the utility construction technique used in the axiomatizations of additive conjoint measurement in, e.g., Krantz *et al.* [19] and Wakker [28].

A few remarks are in order: First, cardinal symmetry holds vacuously for any pair of preferences which do not have essential events. Second, cardinal symmetry does not impose restrictions on the DMs' relative ambiguity attitudes. In fact, for all acts ranked by $\succcurlyeq_i$, the consequence obtained if A_i is always strictly better than that obtained if A_i^c, so that all acts are bets *on the same event* A_i. Intuitively, a DM's ambiguity attitude affects these bets symmetrically, so that his preferences do not convey any information about it. Moreover, cardinal symmetry does not constrain the DMs' relative confidence on A_1 and A_2, since the 'win' (or 'loss') payoffs can be different for the two DMs.

On the other hand, it does unsurprisingly restrict their relative cardinal risk attitudes. To better understand the *relative* restrictions implied by cardinal symmetry, assume that consequences are monetary payoffs and that both DMs like more money to less. Suppose that, when betting on events (A_1, A_2), (9) holds for some 'loss' payoffs x and y and 'win' payoffs $v^* \succ_1 v_*$ and $w^* \succ_2 w_*$ respectively. This says that exchanging v_* for v^* as the prize for A_1, and w_* for w^* as the prize for A_2, can for both DMs be traded off with a reduction in 'loss' from x to y. Suppose that when the initial loss is $x' < x$, $\succcurlyeq_1$ is willing to trade off the increase in 'win' with a reduction in 'loss' to y', but $\succcurlyeq_2$ accepts reducing 'loss' only to $y'' > y'$ (that is, $w_* \, A_2 \, x' \succ_2 w^* \, A_2 \, y'$, in violation of (10)). That is, as the amount of the low payoff decreases, $\succcurlyeq_2$ becomes more sensible to differences in payoffs than $\succcurlyeq_1$. Such diversity of behavior — that we intuitively attribute to differences in the DMs' *risk* attitude — is ruled out by cardinal symmetry, which requires that the two DMs *consistently* agree on the acceptable tradeoff for improving their 'win' payoff, and similarly for the 'loss' payoff. It is important to stress that this discussion makes sense only when both DMs are faced with nontrivial uncertainty (i.e., they are both betting on essential events). Thus, we do not use 'trade-off' to mean certain substitution; rather, substitution in the context of an uncertain prospect.

To see how cardinal symmetry is used to show that two biseparable preferences have the same cardinal risk attitude, assume first that the two relations are **ordinally equivalent**: for every $x, y \in X$, $x \succcurlyeq_1 y \Leftrightarrow x \succcurlyeq_2 y$. When that is the case, cardinal symmetry holds if and only if their canonical utility indices are positive affine transformations of each other. In order to simplify the statements, we write $u_1 \approx u_2$ to denote such 'equality' of indices.

Proposition 6 *Suppose that $\succcurlyeq_1$ and $\succcurlyeq_2$ are ordinally equivalent biseparable preferences which have essential events. Then $\succcurlyeq_1$ and $\succcurlyeq_2$ are cardinally symmetric* if and only if *their canonical utility indices satisfy $u_1 \approx u_2$.*

The intuition of the proof (see Appendix B) can be quickly grasped by rewriting, say, Eqs. (9) and (10) in terms of the canonical representations to find that for every $x, y, x', y' \in X$,

$$u_1(x) - u_1(y) = u_1(x') - u_1(y') \Longleftrightarrow u_2(x) - u_2(y) = u_2(x') - u_2(y').$$

Notice however that this does *not* imply that the preferences are identical on binary acts: The DMs' beliefs on events could be totally different.

The comparative notion of ambiguity aversion we propose in the next subsection checks comparative uncertainty aversion in preferences with the same cardinal risk attitude. Clearly, it would be nicer to have a comparative notion that ranks also preferences without the same cardinal risk attitude. In Subsection 6.1, we discuss how to extend our notion to deal with these cases. This extension requires the exact measurement of the two preferences' canonical utility indices, and is thus 'less behavioral' than the one we just anticipated.

Finally, we remark that a symmetric exercise to that performed here is to coarsen comparative uncertainty aversion so as to rank preferences by their cardinal risk aversion only. In [13] it is shown that for biseparable preferences such ranking is represented by the ordering of canonical utilities by their relative concavity, thus generalizing the standard result.

2.2 Comparative and Absolute Ambiguity Aversion

Having thus prepared the ground, our comparative notion of ambiguity is immediately stated:

Definition 7 *Let $\succcurlyeq_1$ and $\succcurlyeq_2$ be two preference relations. We say that $\succcurlyeq_2$ is* ***more ambiguity averse than*** *$\succcurlyeq_1$ whenever both the following conditions hold:*

(A) $\succcurlyeq_2$ is more uncertainty averse than $\succcurlyeq_1$;

(B) $\succcurlyeq_1$ and $\succcurlyeq_2$ are cardinally symmetric.

Thus, we restrict our attention to pairs which are cardinally symmetric. As explained earlier, when one DM's preference does not have an essential event, cardinal risk aversion does not play a role in that DM's choices, so that we do not need to remove it from the picture.

Remark 8 So far, we have tacitly assumed that cardinal risk and ambiguity attitude completely characterize biseparable preferences. Indeed, the validity of this can be easily verified by observing that if two such preferences are 'as uncertainty averse as' each other (that is, $\succcurlyeq_1$ is more uncertainty averse than $\succcurlyeq_2$, and *vice versa*), they are identical.

We finally come to the absolute definition of ambiguity aversion and love. Let $\geqslant$ be a preference relation on $\mathcal{F}$ with a SEU representation.[5] As we observed in the Introduction, these relations intuitively embody ambiguity neutrality. We propose to use them as the benchmark for defining ambiguity aversion. Of course, one could intuitively hold that the SEU ones are not the *only* relations embodying ambiguity neutrality, and thus prefer using a wider set of benchmarks. This alternative route is discussed in Subsection 6.3 below.

Definition 9 *A preference relation $\succcurlyeq$ is* ***ambiguity averse (loving)*** *if there exists a SEU preference relation $\geqslant$ which is less (more) ambiguity averse than $\succcurlyeq$. It is* ***ambiguity neutral*** *if it is* both *ambiguity averse and ambiguity loving.*

If $\geqslant$ is a SEU preference which is less ambiguity averse than $\succcurlyeq$, we call it a **benchmark preference** for $\succcurlyeq$. We denote by $\mathcal{R}(\succcurlyeq)$ the set of all benchmark preferences for $\succcurlyeq$. That is,

$$\mathcal{R}(\succcurlyeq) \equiv \{\geqslant \subseteq \mathcal{F} \times \mathcal{F} : \geqslant \text{ is SEU and } \succcurlyeq \text{ is more ambiguity averse than } \geqslant\}.$$

Each benchmark preference $\geqslant \in \mathcal{R}(\succcurlyeq)$ induces a probability measure P on Σ, so a natural twin of $\mathcal{R}(\succcurlyeq)$ is the set of the **benchmark measures**:

$$\mathcal{M}(\succcurlyeq) = \{P \in \Delta : P \text{ represents } \geqslant, \text{ for } \geqslant \in \mathcal{R}(\succcurlyeq)\}.$$

Using this notation, Definition 9 can be rewritten as follows: $\succcurlyeq$ is ambiguity averse if either $\mathcal{R}(\succcurlyeq) \neq \emptyset$, or $\mathcal{M}(\succcurlyeq) \neq \emptyset$.

[5] We use the symbols $\geqslant$ (and $>$) to denote SEU weak (and strict) preferences.

3 The Characterizations

We now characterize the notions of comparative and absolute ambiguity aversion defined in the previous section for the general case of biseparable preferences, and the important subcases of CEU and MEU preferences. To start, we use Proposition 6 and the observation that the canonical utility index of a preference with no essential events is ordinal, to show that if two preferences are biseparable and they are ranked by Definition 7, they have the same canonical utility index:

Theorem 10 *Suppose that $\succcurlyeq_1$ and $\succcurlyeq_2$ are biseparable preferences, and that $\succcurlyeq_2$ is more ambiguity averse than $\succcurlyeq_1$. Then $u_1 \approx u_2$.*

Checking cardinal symmetry is clearly not a trivial task, but for an important subclass of preference relations — the c-linearly biseparable preferences in an Anscombe-Aumann setting — it is implied by comparative uncertainty aversion. In fact, under c-linearity, ordinal equivalence easily implies cardinal symmetry, so that we get:

Proposition 11 *Suppose that X is a convex subset of a vector space, and that $\succcurlyeq_1$ and $\succcurlyeq_2$ are c-linearly biseparable preferences. $\succcurlyeq_2$* is more ambiguity averse than *$\succcurlyeq_1$ if and only if $\succcurlyeq_2$* is more uncertainty averse than *$\succcurlyeq_1$.*

Therefore, in this case Definition 4 can be *directly* used as our definition of comparative ambiguity attitude.

3.1 Absolute Ambiguity Aversion

We first characterize absolute ambiguity aversion for a general biseparable preference $\succcurlyeq$. Suppose that V is a canonical representation of $\succcurlyeq$, with canonical utility u. We let

$$\mathcal{D}(\succcurlyeq) \equiv \left\{ P \in \Delta : \int_S u(f(s))\, P(ds) \geq V(f) \text{ for all } f \in \mathcal{F} \right\}.$$

That is, $\mathcal{D}(\succcurlyeq)$, which depends only on V, is the set of beliefs inducing preferences which assign (weakly) higher expected utility to every act f. These preferences exhaust the set of the benchmarks of $\succcurlyeq$:

Theorem 12 *Let $\succcurlyeq$ be a biseparable preference. Then, $\mathcal{M}(\succcurlyeq) = \mathcal{D}(\succcurlyeq)$. In particular, $\succcurlyeq$ is ambiguity averse* if and only if *$\mathcal{D}(\succcurlyeq) \neq \emptyset$.*

Let ρ be the capacity associated with the canonical representation V. It is immediate to see that if $P \in \mathcal{D}(\succcurlyeq)$, then $P \geq \rho$. Thus, non-emptiness of the core of ρ (the set of the probabilities that dominate ρ pointwise, that we denote $\mathcal{C}(\rho)$) is necessary for $\succcurlyeq$ to be ambiguity averse. In Subsection 3.2 it is shown to be not sufficient in general.

Turn now to the characterization of ambiguity aversion for the popular CEU and MEU models. Suppose first that $\succcurlyeq$ is a CEU preference relation represented by the capacity ν, and let $\mathcal{C}(\nu)$ denote ν's possibly empty core. It is shown that $\mathcal{D}(\succcurlyeq) = \mathcal{C}(\nu)$, so that the following result — which also provides a novel decision-theoretic interpretation of the core as the set of all the benchmark measures — follows as a corollary of Theorem 12.

Corollary 13 *Suppose that $\succcurlyeq$ is a CEU preference relation, represented by capacity ν. Then $\mathcal{C}(\nu) = \mathcal{M}(\succcurlyeq)$. In particular, $\succcurlyeq$ is ambiguity averse* if and only if *$\mathcal{C}(\nu) \neq \emptyset$.*

Thus, the core of an ambiguity averse capacity is *equal* to the set of its benchmark measures, and the ambiguity averse capacities are those with a non-empty core, called 'balanced'. A classical result (see, e.g., Kannai [17]) thus provides an internal characterization of ambiguity aversion in the CEU case: Letting 1_A denote the characteristic function of $A \in \Sigma$, a capacity reflects ambiguity aversion *if and only if* for all $\lambda_1, \ldots, \lambda_n \geq 0$ and all $A_1, \ldots, A_n \in \Sigma$ such that $\sum_{i=1}^n \lambda_i 1_{A_i} \leq 1_S$, we have $\sum_{i=1}^n \lambda_i \nu(A_i) \leq 1$. As convex capacities are balanced, but not conversely, the corollary motivates our claim that convexity does not characterize our notion of ambiguity aversion. This point is illustrated by Example 25 below, which presents a capacity that intuitively reflects ambiguity aversion but is not convex.

On the other hand, given a MEU preference relation $\succcurlyeq$ with set of priors C, it is shown that $\mathcal{D}(\succcurlyeq) = C$. Thus, Theorem 12 implies that any MEU preference is ambiguity averse (as it is intuitive) and, more interestingly, that the set C can be interpreted as the set of the benchmark measures for $\succcurlyeq$.

Corollary 14 *Suppose that $\succcurlyeq$ is a MEU preference relation, represented by the set of probabilities C. Then $C = \mathcal{M}(\succcurlyeq)$, so that $\succcurlyeq$ is ambiguity averse.*

As to ambiguity love, reversing the proof of Theorem 12 shows that for any biseparable preference, ambiguity love is characterized by nonemptiness of the set

$$\mathcal{E}(\succcurlyeq) \equiv \left\{ P \in \Delta : \int_S u(f(s))\, P(ds) \leq V(f) \text{ for all } f \in \mathcal{F} \right\}.$$

In particular, a CEU preference with capacity ν is ambiguity loving *if and only if* the set of probabilities dominated by ν is non-empty. As for MEU preferences: None is ambiguity loving. Conversely, any 'maximax' EU preference is ambiguity loving, with $\mathcal{E}(\succcurlyeq) = C$.

Finally, we look at ambiguity neutrality. Since we started with an *informal* intuition of SEU preferences as reflecting neutrality to ambiguity, an important consistency check on our analysis is to verify that they are ambiguity neutral in the *formal* sense. This is the case:

Proposition 15 *Let $\succcurlyeq$ be a biseparable preference. Then $\succcurlyeq$ is ambiguity neutral* if and only if *it is a SEU preference relation.*

3.2 Comparative Ambiguity Aversion

We conclude the section with the characterization of comparative ambiguity aversion. The general result on comparative ambiguity, an immediate consequence of Theorem 12, is stated as follows (where ρ_1 and ρ_2 represent the willingness to bet of $\succcurlyeq_1$ and $\succcurlyeq_2$ respectively):

Proposition 16 *Let $\succcurlyeq_1$ and $\succcurlyeq_2$ be two biseparable preferences.* If *$\succcurlyeq_2$ is more ambiguity averse than $\succcurlyeq_1$,* then *$\rho_1 \geq \rho_2$, $\mathcal{D}(\succcurlyeq_1) \subseteq \mathcal{D}(\succcurlyeq_2)$, $\mathcal{E}(\succcurlyeq_1) \supseteq \mathcal{E}(\succcurlyeq_2)$ and $u_1 \approx u_2$.*

Thus, relative ambiguity implies containment of the sets $\mathcal{D}(\succcurlyeq)$ and $\mathcal{E}(\succcurlyeq)$ (clearly in opposite directions), and dominance of the willingness to bet ρ. Of course, the proposition lacks a converse, and thus it does not offer a full characterization. As we argue below, biseparable preferences seem to have too little structure for obtaining a general characterization result.

Things are different if we restrict our attention to specific models. For instance, the next result characterizes comparative ambiguity for the CEU and MEU models:

Theorem 17 *Let $\succcurlyeq_1$ and $\succcurlyeq_2$ be biseparable preferences, with canonical utilities u_1 and u_2 respectively.*

(*i*) *Suppose that $\succcurlyeq_1$ and $\succcurlyeq_2$ are CEU, with respective capacities ν_1 and ν_2. Then $\succcurlyeq_2$ is more ambiguity averse than $\succcurlyeq_1$* if and only if *$\nu_1 \geq \nu_2$ and $u_1 \approx u_2$.*

(*ii*) *Suppose that $\succcurlyeq_1$ is MEU, with set of probabilities C_1. Then $\succcurlyeq_2$ is more ambiguity averse than $\succcurlyeq_1$* if and only if *$C_1 = \mathcal{D}(\succcurlyeq_1) \subseteq \mathcal{D}(\succcurlyeq_2)$ and $u_1 \approx u_2$.*

Observe that part (*ii*) of the proposition does more than characterize comparative ambiguity for MEU preferences, as it applies to any biseparable $\succcurlyeq_2$. For instance, it is immediate to notice that one can characterize *absolute* ambiguity aversion using that result and the fact that if $\succcurlyeq_1$ is a SEU preference relation with beliefs P, then $C_1 = \{P\}$. Also, a symmetric result to (*ii*) holds: If $\succcurlyeq_2$ is 'maximax' EU, it is more ambiguity averse than $\succcurlyeq_1$ iff $C_2 = \mathcal{E}(\succcurlyeq_2) \subseteq \mathcal{E}(\succcurlyeq_1)$.

Remark 18 Proposition 17 can be used to explain the apparent incongruence of the characterization of comparative risk aversion in SEU (in the sense of Yaari [30]) and of comparative ambiguity aversion in CEU: Convexity of ν *seems* to be the natural counterpart of concavity of u, but it is not. This is due to the different uniqueness properties of utility functions and capacities. A SEU $\geqslant_2$ is more risk averse than a SEU $\geqslant_1$ iff for every common normalization of the utilities, we have $u_2(x) \geq u(x)$ inside the interval of normalization. Since any normalization is allowed, u_2 must then be a concave transformation of u_1. In the case of capacities only one normalization is allowed, so we only have $\nu_1 \geq \nu_2$.

It is not difficult to show that the necessary conditions of Proposition 16 are not sufficient if taken one by one. For instance, there are pairs of MEU (resp. CEU) preferences $\succcurlyeq_1$ and $\succcurlyeq_2$ such that $\rho_1 \geq \rho_2$ (resp. $\mathcal{C}(\nu_1) = \mathcal{D}(\succcurlyeq_1) \subseteq \mathcal{D}(\succcurlyeq_2) = \mathcal{C}(\nu_2)$) does not entail that $\succcurlyeq_2$ is more ambiguity averse than $\succcurlyeq_1$.

Example 19 Let $S = \{s_1, s_2, s_3\}$, Σ the power set of S. Consider the probabilities P, Q and R defined by $P = [1/2, 0, 1/2]$, $Q = [0, 1, 0]$ and $R = [1/2, 1/2, 0]$. Let C_1 and C_2 respectively be the closed convex hull of $\{P, Q\}$ and $\{P, Q, R\}$. Then, $\rho_1 = \underline{P}_1 = \underline{P}_2 = \rho_2$, but $C_2 \not\subseteq C_1$, and indeed by Prop. 17 the MEU preference $\succcurlyeq_2$ inducing C_2 is more ambiguity averse than the MEU preference $\succcurlyeq_1$ inducing C_1.

Consider next a capacity ν such that $\nu(A) = 1/3$ for any $A \neq \{\emptyset, S\}$, and a probability P' equal to 1/3 on each singleton. Then $\mathcal{C}(\nu) = \{P'\}$, so that ν is balanced, but not exact (for instance, $P'(\{s_1, s_2\}) = 2/3 > 1/3 = \nu(\{s_1, s_2\})$). We have $\mathcal{C}(\nu) \subseteq \mathcal{C}(P')$ but $\nu \not\geq P'$, and by Prop. 17 the CEU preference inducing ν is *not* more ambiguity averse than that inducing P'. In contrast, P' is exact, and we have both $\mathcal{C}(P') \subseteq \mathcal{C}(\nu)$ and $P' \geq \nu$. △

These examples illustrate two conceptual observations. The first (anticipated in Subsection 3.1) is that non-emptiness of the core of ρ is not sufficient for absolute ambiguity aversion: A probability can dominate ρ without being a benchmark measure for $\succcurlyeq$. Unsurprisingly, in general the capacity ρ does not completely describe the DM's ambiguity attitude. The second observation is that, while $\mathcal{D}(\succcurlyeq)$ *does* characterize the DM's absolute ambiguity aversion, it is also an incomplete description of the DM's ambiguity attitude: There can be preferences $\succcurlyeq_1$ and $\succcurlyeq_2$ strictly ranked by comparative ambiguity even though $\mathcal{D}(\succcurlyeq_1) = \mathcal{D}(\succcurlyeq_2)$.

To better appreciate the difficulty of obtaining a general sufficiency result for biseparable preferences, we now present an example in which *all* the necessary conditions hold but the comparative ranking does not obtain.

Example 20 For a general S and Σ (but see the restriction on P below), consider two preference relations $\succcurlyeq_1$ and $\succcurlyeq_2$ which behave according to example (*iii*) of biseparable preference in Section 1. Both have identical P and u (which ranges in a nondegenerate interval of $\mathbb{R}$), with the following restriction on P: There are at least three disjoint events in Σ, A_1, A_2 and A_3 such that $P(A_i) > 0$ for $i = 1, 2, 3$ (otherwise both preferences are indistinguishable from SEU preferences with utility u and beliefs P). Their β parameters are different, in particular $\beta_2 > \beta_1 > 0$. Clearly $\rho_1 = \rho_2 = P$ and $u_1 \approx u_2$. It is also immediate to verify that, under the assumption on P, $\mathcal{D}(\succcurlyeq_1) = \mathcal{D}(\succcurlyeq_2) = \{P\}$ and $\mathcal{E}(\succcurlyeq_1) = \mathcal{E}(\succcurlyeq_2) = \emptyset$, so that both preferences are (strictly) ambiguity averse. However, $\succcurlyeq_1$ is *not* more ambiguity averse than $\succcurlyeq_2$ (nor are $\succcurlyeq_1$ and $\succcurlyeq_2$ equal, which would follow from two applications of the converse). Indeed, the parameter β measures comparative ambiguity for these preferences, so that $\succcurlyeq_2$ is more ambiguity averse than $\succcurlyeq_1$. △

4 Unambiguous Acts and Events

Let $\succcurlyeq$ be an ambiguity averse or loving preference relation. Even though the preference relation has a strict ambiguity attitude, it may nevertheless behave in an ambiguity neutral fashion with respect to some subclass of acts and events, that we may like to consider 'unambiguous'. The purpose of this section is to identify the class of the unambiguous acts and the related class of unambiguous events, and to present a characterization of the latter for biseparable (in particular CEU and MEU) preference relations. We henceforth focus on ambiguity averse preference relations, but it is easy to see that all the results in this section can be shown for ambiguity *loving* preferences. A more extensive discussion of the behavioral definition of ambiguity for events and acts is found in our [14].

In view of our results so far, the natural approach in defining the class of unambiguous events of a preference relation $\succcurlyeq$ is to fix a a benchmark $\geqslant \in \mathcal{R}(\succcurlyeq)$, and to consider the subset of all the acts in $\mathcal{F}$ over which $\succcurlyeq$ is as ambiguity averse as $\geqslant$. Intuitively, ambiguity is a property that the DM attaches to partitions of events, so that nonconstant acts which generate the same partition should be consistently deemed either both ambiguous or both unambiguous. Hence, we consider as 'truly' unambiguous only the acts which belong to the set defined below.

Definition 21 *Given a preference relation $\succcurlyeq$ and $\geqslant \in \mathcal{R}(\succcurlyeq)$, the set of **$\geqslant$-unambiguous acts**, denoted $\mathcal{H}_{\geqslant}$, is the largest subset of $\mathcal{F}$ satisfying the following two conditions:*[6]

(A) For every $x \in X$ and every $f \in \mathcal{H}_{\geqslant}$, $\succcurlyeq$ and $\geqslant$ agree on the ranking of f and x.

(B) For every $f \in \mathcal{H}_{\geqslant}$ and every $g \in \mathcal{F}$, if $\{g^{-1}(\{x\}) : x \in X\} \subseteq \{f^{-1}(\{x\}) : x \in X\}$, then $g \in \mathcal{H}_{\geqslant}$.

Given a preference relation $\succcurlyeq$, for any $f \in \mathcal{F}$ denote by Γ_f the collection of all the 'upper pre-image' sets of f, that is,

$$\Gamma_f = \{\{s : f(s) \succcurlyeq x\} : x \in X\}. \tag{13}$$

Since any benchmark $\geqslant \in \mathcal{R}(\succcurlyeq)$ is ordinally equivalent to $\succcurlyeq$, for any act $f \in \mathcal{F}$ the upper pre-images of f with respect to $\succcurlyeq$ and $\geqslant$ coincide: for all $x \in X$, $\{s : f(s) \succcurlyeq x\} = \{s : f(s) \geqslant x\}$. The set $\Lambda_{\geqslant} \subseteq \Sigma$ of the **$\geqslant$-unambiguous events** is thus naturally defined to be the collection of all sets of upper pre-images of the acts in $\mathcal{H}_{\geqslant}$. That is,

$$\Lambda_{\geqslant} \equiv \bigcup_{f \in \mathcal{H}_{\geqslant}} \Gamma_f.$$

It is immediate to observe that if $A \in \Lambda_{\geqslant}$, then for every $x, y \in X$ the binary act $x\,A\,y$ belongs to $\mathcal{H}_{\geqslant}$. This implies that $A^c \in \Lambda_{\geqslant}$ (that is, $\Lambda_{\geqslant}$ is closed w.r.t. complements).

[6] *Such set is well-defined since it is trivially true that the union of any collection of sets satisfying* (A) *and* (B) *below also satisfies the two conditions.*

We now present the characterization of the set $\Lambda_{\geqslant}$. This turns out to be quite simple and intuitive: It is the subset of the events over which the capacity ρ representing $\succcurlyeq$'s willingness to bet is **complement-additive** (sometimes called 'symmetric'):

Proposition 22 *Let $\succcurlyeq$ be an ambiguity averse biseparable preference with willingness to bet ρ. Then for every $\geqslant \in \mathcal{R}(\succcurlyeq)$, the set $\Lambda_{\geqslant}$ satisfies:*

$$\Lambda_{\geqslant} = \{A \in \Sigma : \rho(A) + \rho(A^c) = 1\}. \tag{14}$$

It immediately follows from the proposition that the choice of the specific benchmark $\geqslant$ does not change the resulting set of events. In light of this, we henceforth call $\Lambda = \Lambda_{\geqslant}$ the **set of unambiguous events for $\succcurlyeq$**.

The consequences of the proposition for the CEU and MEU models are clear: Just substitute ν or $\underline{P}$ for ρ. In particular, when $\succcurlyeq$ is a MEU preference with a set of probabilities C, it can be further shown that Λ is the set of events on which all probabilities agree:

$$\Lambda = \{A \in \Sigma : \rho(A) = P(A) \text{ for all } P \in C\}.$$

It is also interesting to observe that Λ is in general not an algebra. This is intuitive, as the intersection of unambiguous events could be ambiguous.[7]

As to the set of unambiguous acts $\mathcal{H}_{\geqslant}$, it can also be seen to be independent of the choice of benchmark. In general, the only way to ascertain which acts are unambiguous is to construct the set $\mathcal{H}_{\geqslant}$. However, for MEU preferences and for CEU preferences whose capacity is exact (the lower envelope of its core), the set $\mathcal{H}_{\geqslant}$ is the set of all the acts which are measurable with respect to the events in Λ. Therefore, in these cases Λ characterizes the set of unambiguous acts as well. (All these results are proved in [14].)

5 Back to Ellsberg

We now illustrate our results using the classical Ellsberg urn. The urn contains 90 balls of three colors: red, blue and yellow. The DM knows that there are 30 red balls and that the other 60 balls are either blue or yellow. However, he does not know their relative proportion. The state space for an extraction from the urn is $S = \{B, R, Y\}$. Given the nature of his information, it is natural to assume that the DM's preference relation $\succcurlyeq$ will be such that its set of unambiguous events satisfies $\Lambda \supseteq \{\emptyset, \{R\}, \{B, Y\}, S\}$.

In particular, assume that the DM's preference relation is CEU and it induces the capacity ν. To reflect the fact that $\{R\}$ and $\{B, Y\}$ form an unambiguous partition, we

[7] See Zhang [31] for a compelling urn example in which this happens.

know from the previous section that if the DM is ambiguity averse (or loving) ν must satisfy

$$\nu(R) + \nu(B, Y) = 1. \tag{15}$$

Also, because of the symmetry of the information that the DM is given, it is natural to assume that

$$\nu(B) = \nu(Y) \quad \text{and} \quad \nu(B, R) = \nu(R, Y). \tag{16}$$

We first show that, if the ambiguity restriction (15) is imposed, ambiguity aversion is not compatible with the following beliefs, which induce behavior that would *on intuitive grounds* be considered 'ambiguity loving':

$$\begin{array}{ccccc} \nu(R) & < & \nu(B) & = & \nu(Y); \\ \nu(B, Y) & < & \nu(B, R) & = & \nu(R, Y). \end{array} \tag{17}$$

Proposition 23 *No ambiguity averse CEU preference relation such that its set of unambiguous events contains $\{\{R\}, \{B, Y\}\}$ can agree with the ranking (17).*

In his paper on ambiguity aversion, Epstein [8] also discusses the Ellsberg urn, and he presents a convex capacity compatible with ambiguity loving in his sense (see Subsection 6.3 for a brief review), which satisfies the conditions in (17). This is the capacity ν_1 defined by

$$\nu_1(R) = \frac{1}{12}, \quad \nu_1(B, Y) = \frac{1}{3},$$

$$\nu_1(B) = \nu_1(Y) = \frac{1}{6}, \quad \nu_1(B, R) = \nu_1(R, Y) = \frac{1}{2}.$$

He thus concludes that convexity of beliefs does not imply ambiguity aversion for CEU preferences (it is also not implied, in his definition).

We know from Corollary 13 that convexity implies ambiguity aversion in our sense. Proposition 23 helps clarifying why this example does not conflict with the intuition developed earlier: In fact, ν_1 does embody ambiguity aversion in our sense, but it does not reflect the usual presumption that $\{R\}$ and $\{B, Y\}$ are seen as unambiguous events. If it did, it would have to satisfy (15), which is not the case (it cannot be, since convex capacities are balanced). For us, the DM with beliefs ν_1 does not perceive $\{R\}$ and $\{B, Y\}$ as unambiguous. Of course, then it is not clear in which sense the conditions in (17) should 'intuitively' embody ambiguity loving behavior.

Going back to the example, we would say that the DM's preferences *intuitively* reflect ambiguity aversion if the reverse inequalities held:

$$\begin{array}{ccccc} \nu(R) & \geq & \nu(B) & = & \nu(Y); \\ \nu(B, Y) & \geq & \nu(B, R) & = & \nu(R, Y). \end{array} \tag{18}$$

We now show that the notion of ambiguity aversion proposed earlier characterizes this intuitive ranking when, besides the obvious symmetry restrictions in (16), we strengthen the requirement in (15) in the following natural way:

$$\nu(R) = \frac{1}{3} \quad \text{and} \quad \nu(B, Y) = \frac{2}{3}. \tag{19}$$

Proposition 24 *Let $\succcurlyeq$ be a CEU preference relation such that its representing capacity ν satisfies the equalities (16) and (19). Then $\succcurlyeq$ is ambiguity averse* if and only if *ν agrees with the ranking (18).*

In closing our discussion of Ellsberg's problem, we provide further backing for our belief that convexity is not necessary for ambiguity aversion. Here is a capacity which is not convex, and still makes the typical Ellsberg choices.

Example 25 Consider the capacity ν_2 defined by (19) and

$$\nu_2(B) = \nu_2(Y) = \frac{7}{24}, \quad \nu_2(B, R) = \nu_2(R, Y) = \frac{1}{2}.$$

This capacity satisfies (18), so that it reflects ambiguity aversion both formally and intuitively, but it is not superadditive, let alone convex. △

6 Discussion

In this section we discuss some of the choices we have made in the previous sections. First we briefly discuss how the comparative ambiguity ranking can be extended to preferences with different cardinal risk attitude. Then we discuss in more detail how the unambiguous acts described in Section 4 can be used in the comparative ranking, and why we chose SEU preferences as benchmarks.

6.1 Comparative Ambiguity and Equality of Cardinal Risk Attitude

As we observed earlier, our comparative ambiguity aversion notion cannot compare biseparable preferences with different canonical utility indices. Of course, the characterization results of Section 3 can be used to *qualitatively* compare two preference by ambiguity: For instance, we can look at two CEU preferences and compare their willingness to bet, or we can use utility functions to compare two SEU preferences by risk aversion, even if they do not have the same beliefs.

However, when dealing with biseparable preferences, it is easy to apply the intuition of our comparative ranking to compare preferences which do not have the same canonical utility. This requires eliciting the canonical utility indices first, and then using acts and constants that are 'utility equivalents' in Eqs. (7) and (8).[8] The ranking thus obtained is very general (it does not even entail ordinal equivalence), but it yields *mutatis mutandis* the same characterization results that we obtained with the more restrictive one. For instance: $\succcurlyeq$ is ambiguity averse iff $\mathcal{D}(\succcurlyeq) \neq \emptyset$, and CEU (MEU) preference $\succcurlyeq_2$ is more ambiguity averse than CEU (MEU) preference $\succcurlyeq_1$ iff $\nu_1 \geq \nu_2$ ($C_1 \subseteq C_2$) (but of course in general $u_1 \not\approx u_2$). Nonetheless, this ranking requires the full elicitation of the DMs' canonical utility indices, and is thus operationally more complex than that in Definition 7.

6.2 Using Unambiguous Acts in the Comparative Ranking

One of the intuitive assumptions that our analysis builds on is that constant acts are primitively 'unambiguous': That is, we *assume* that every DM perceives constants as unambiguous. No other acts are 'unambiguous' in this primitive sense. However, one could argue that it is natural to use in the comparative ranking also those acts which are *revealed* to be deemed unambiguous by both DMs, even if they are not constant.

Suppose that $\succcurlyeq$ is an ambiguity averse biseparable preference, and let $\mathcal{H}_{\geqslant}$ ($\Lambda_{\geqslant}$) be its set of unambiguous acts (events), as defined in Section 4. It is possible to see [14] that for every $\geqslant \in \mathcal{R}(\succcurlyeq)$ and every $h \in \mathcal{H}_{\geqslant}$ and $f \in \mathcal{F}$, we have

$$h \geqslant f \Rightarrow h \succcurlyeq f \qquad \text{and} \qquad h > f \Rightarrow h \succ f. \tag{20}$$

That is, all benchmarks according to Definition 7 satisfy the stronger comparative ranking suggested above. Conversely, it is obvious that if $\succcurlyeq$ and a SEU preference $\geqslant$ are cardinally symmetric and satisfy (20), they satisfy Definition 7. Thus, modifying Definition 7 to have (20) in part (A) *does not* change the set of the ambiguity averse preferences.

6.3 A More General Benchmark

We chose SEU maximization as the benchmark representing ambiguity neutrality. While few would disagree that SEU preferences are 'ambiguity neutral' (in a primitive, non-formal sense), some readers may find that the result of Proposition 15 that SEU maximization characterizes ambiguity neutrality does not agree with their intuition of what constitutes ambiguity neutral behavior. In particular, they might feel that we should also classify as ambiguity neutral any non-SEU preference whose likelihood relation can still be represented by a probability measure. This would clearly be the case if we let such preferences be benchmarks for our comparative ambiguity notion. Here we explain why we have not followed that route, and the consequences of this choice for the interpretation of our notions.

[8] For *any* pair of biseparable preferences which have essential events, this elicitation can be done without extraneous devices by using the tradeoff method briefly outlined in Appendix B.

The non-SEU preferences in question are those that are *probabilistically sophisticated* (PS) in the sense of Machina and Schmeidler [20]. For example, consider a CEU preference $\succcurlyeq$ whose willingness to bet is $\rho = g(P)$ for some probability measure P and 'distortion' function g; that is, an increasing $g : [0,1] \to [0,1]$ such that $g(0) = 0$ and $g(1) = 1$. Such $\succcurlyeq$ is PS since its ranking of bets (likelihood relation) is represented by the probability P, but it is not SEU if g is different from the identity function. According to the point of view suggested above, such $\succcurlyeq$ is 'ambiguity neutral'; it should thus be used as a benchmark in characterizing ambiguity aversion. Moreover, if we used PS preferences as benchmarks it might be possible to avoid attributing to ambiguity aversion the effects of probabilistic risk aversion. However, go back to the ambiguous urn of Example 1 and consider the following:

Example 1 (continued) In the framework of Example 1, consider a third DM with CEU preferences $\succcurlyeq_3$, with canonical utility $u(x) = x$ and willingness to bet defined by

$$\rho_3(B) = \frac{1}{4} \quad \text{and} \quad \rho_3(R) = \frac{1}{4}.$$

It is immediate to verify that according to Definition 7, DM 3 is more ambiguity averse than DM 1 (who is SEU), so that he is ambiguity averse in our sense. That seems quite natural, since he is willing to invest less in bets on the ball extractions. With PS benchmarks, we conclude that *both* DMs are ambiguity neutral, since their willingness to bet are ordinally equivalent to the probability ρ_1 ($\rho_3 = g(\rho_1)$ for any distortion g such that $g(1/2) = 1/4$), so that both are PS. Hence, DM 3's behavior is *only* due to his probabilistic risk aversion. Yet, it seems that the fact that DM 3 is only willing to bet 1/4 utils on any color may *at least in part* be due to the ambiguity of the urn and his possible ambiguity aversion. △

This example is not the only case in which using PS benchmarks yields counterintuitive conclusions. When the state space is finite, if we use PS preferences as benchmarks we find that almost every CEU preference inducing a strictly positive ρ on a finite state space is both ambiguity averse and loving. Thus, a large set of preferences are *shown to be* ambiguity neutral. Including, as the following example illustrates, many preferences which are *not* PS.

Example 26 Suppose that two DMs are faced with the following decision problem. There are two urns, both containing 100 balls, either red or black. The DMs are told that Urn I contains at least 40 balls of each color, while Urn II contains at least 10 balls of each color. One ball will be extracted from each urn. Thus, the state space is $S = \{Rr, Rb, Br, Bb\}$, where the upper (lower) case letter stands for the color of the ball extracted from Urn I (II). Suppose that both DMs have CEU preferences $\succcurlyeq_1$ and $\succcurlyeq_2$, with respective willingness to bet ρ_1 and ρ_2. Using obvious notation, suppose that $\rho_1(b) = \rho_1(r) = 0.1$ and $\rho_1(B) = \rho_1(R) = 0.4$, that $\rho_1(s) = 0.04$ for each singleton s, and for every other event ρ_1 is obtained by additivity. According to Definition 9, DM 1 is strictly ambiguity averse. In contrast, with PS benchmarks the result mentioned above shows that he is ambiguity neutral.

Let ρ_2 be as follows: $\rho_2(b) = \rho_2(r) = 0.9$ and $\rho_2(B) = \rho_2(R) = 0.6$, $\rho_2(s) = 0.54$ for each singleton s, $\rho_2(A) = 0.92$ for each $A \in \{Rr \cup Bb, Rb \cup Br\}$, and $\rho_2(A) = 0.95$ for each ternary set. According to Definition 9, DM 2 is ambiguity loving, but if we use PS benchmarks we conclude that she is ambiguity neutral. Both conclusions go against our intuition. Moreover, since both ρ_1 and ρ_2 are not ordinally equivalent to a probability, $\succcurlyeq_1$ and $\succcurlyeq_2$ are not PS. $\triangle$

The foregoing discussion shows some of the difficulties that may arise if we use PS, rather than SEU, preferences as benchmarks with *our* comparative ambiguity aversion notion: We end up attributing too much explanatory power to probabilistic risk aversion. Instead, with SEU benchmarks we overemphasize the role of ambiguity aversion. Is it possible to remove probabilistic risk attitude from the picture, as we did for cardinal risk attitude?[9]

6.3.1 Removing Probabilistic Risk Aversion

Suppose that there is a subset $\mathcal{E}$ of acts which are universally accepted as 'unambiguous', in the sense that we are *sure* that a DM's choices among these acts are unaffected by his ambiguity attitude. Then, if $\mathcal{E}$ (and the associated set of 'unambiguous' events, denoted Γ) is sufficiently rich, we can discriminate between probabilistic risk and ambiguity aversion. For instance, modify Example 1 by assuming the availability of an 'unambiguous' randomizing device, so that each state describes the result of the device as well. Now, find a set A of results of the device (obviously, here Γ is the family of all such sets) which is as likely as R(ed) and then check if B(lack) is as likely as A^c. If it is, the DM behaves identically when faced with (equally likely) ambiguous and unambiguous events, so that all the non-additivity of ρ_3 on $\{B, R\}$ must be due to his probabilistic risk aversion. His preferences are also PS on the extended problem. If it is not, then DM 3's behavior is affected by ambiguity, and his preferences are not PS on the extended problem. The point is that in the presence of a sufficiently rich Γ, a DM whose preferences are PS is treating ambiguous and unambiguous events symmetrically, and is hence *intuitively* ambiguity neutral. Therefore, in such a case we would expect PS preferences to be found ambiguity neutral. This is not the case in the original version of Example 1, since a rich set of 'unambiguous' events is missing.

More generally, consider a biseparable preference $\succcurlyeq$ which is not PS overall, but is PS when comparing *only* unambiguous acts. That is, the DM behaves as if he forms a probability P on the set Γ, and calculates his willingness to bet on these events by means of a distortion function g which only reflects his probabilistic risk attitude. As we did in controlling for cardinal risk attitude, we want to use as benchmarks for $\succcurlyeq$ only those PS preferences — that with a small abuse of notation we also denote $\succcurlyeq$ — which

[9] We thank Peter Klibanoff for his substantial help in developing the ensuing discussion.

have the same probabilistic risk attitude; e.g., those biseparable preferences which share g as distortion function. Interestingly, it turns out that if the set $\mathcal{E}$ is rich enough, any PS preference $\succcurlyeq$ satisfying Eq. (20) for all $h \in \mathcal{E}$ has this property. This is exactly the approach followed by Epstein [8] in his work on ambiguity aversion: He assumes the existence of a suitably rich set Γ of 'unambiguous' events,[10] defines $\mathcal{E}$ as the set of all the Γ-measurable acts, and uses Eq. (20) with $h \in \mathcal{E}$ as his comparative ambiguity notion. His choice of benchmark are PS preferences.

This approach attains the objective of 'filtering' the effects of probabilistic risk attitude from our absolute ambiguity notion. It thus yields a finer assessment of the DM's ambiguity attitude. However, the foregoing discussion has illustrated that a crucial ingredient to this filtration is the existence of a set of 'unambiguous' acts which is sufficiently rich: If it is too poor (e.g., it contains only the constants, as in Example 26), we may use benchmarks whose probabilistic risk attitude is different from the DM's. This may cause Epstein's approach to reach counterintuitive conclusions, as illustrated in the previous examples.

The main problem we have with this approach is that we find it undesirable to base our measurement of ambiguity attitude on an exogenous notion of 'ambiguity', especially in view of the richness requisite. It seems that in many cases of interest the 'obvious' set of 'unambiguous' acts does not satisfy such requisite; e.g., Ellsberg's example. Our objective is to develop a notion of ambiguity attitude which is based on the weakest set of primitive requisites (like the two assumptions stated in the Introduction), even though this has a cost in terms of the 'purity' of the interpretation of the behavioral feature we measure.

Epstein and Zhang [10] propose a behavioral foundation to the notion of 'ambiguity', so that the existence of a rich set $\mathcal{E}$ can be *objectively* verified, solving the problem mentioned above. In [14] we present an example which suggests that their behavioral notion can lead to counterintuitive conclusions (in that case, an intuitively ambiguous event is found unambiguous). More generally, we see the following problem with this enterprise: There may be events which are 'unambiguous' (resp. 'ambiguous') with respect to which the DM nonetheless behaves in an ambiguity non-neutral (resp. neutral) fashion. Consider a DM who listens to a weather forecast stated as a probabilistic judgement. If the DM does not consider the specific source reliable, he might express a willingness to bet which is a distortion of this judgement, while being probabilistic risk neutral. Alternatively, he may find the source reliable, hence perceive no ambiguity, but be probabilistically risk averse. A preference-based notion of ambiguity must be able to distinguish between these two cases, classifying the relevant events ambiguous in the first case and unambiguous in the second. And this without using any auxiliary information. Considering moreover that the set of 'verifiably unambiguous' events must be rich, we are skeptical that this feat is possible: The problem is that the Savage set-up does not

[10] The richness condition is: For every $F \subseteq E$ in Σ and $A \in \Gamma$ such that A is as likely as E, there is $B \subseteq A$ in Γ such that B is as likely as F. Epstein remarks that richness of Γ is not required for some of his results.

provide us with enough instruments; it is too abstract.

6.3.2 Summing Up

We have argued that what motivates using PS (rather than SEU) preferences as benchmarks is the objective of discriminating between probabilistic risk aversion and ambiguity attitude. We have shown that this requires a rich set of 'verifiably unambiguous' events, and briefly reviewed our doubts about the possibility of providing a behavioral foundation to this 'verifiable ambiguity' notion in a general subjective setting without extraneous devices. In contrast, the analysis in this paper shows that there are no such problems in using SEU benchmarks to identify an 'extended' notion of ambiguity attitude, which can be disentangled from cardinal risk attitude using only behavioral data and no extraneous devices. Though it does not distinguish 'real' ambiguity and probabilistic risk attitudes, we think that this 'extended' ambiguity attitude is worthwhile, especially because of its wider applicability.

Appendix A Capacities and Choquet Integrals

A set-function ν on (S, Σ) is called a **capacity** if it is monotone and normalized. That is: if for $A, B \in \Sigma$, $A \subseteq B$, then $\nu(A) \leq \nu(B)$; $\nu(\emptyset) = 0$ and $\nu(S) = 1$. A capacity is called a **probability measure** if it is **finitely additive**: $\nu(A \cup B) = \nu(A) + \nu(B)$ for all A disjoint from B. It is called **convex** if for every pair $A, B \in \Sigma$, we have $\nu(A \cup B) \geq \nu(A) + \nu(B) - \nu(A \cap B)$.

The **core** of a capacity ν is the (possibly empty) set $\mathcal{C}(\nu)$ of all the probability measures on (S, Σ) which dominate it, that is,

$$\mathcal{C}(\nu) \equiv \{P : P \in \Delta,\ P(A) \geq \nu(A) \text{ for all } A \in \Sigma\}.$$

Following the usage in Cooperative Game Theory (e.g., Kannai [17]), all capacities with nonempty core are called **balanced**. A capacity ν is called **exact** it is balanced and it is equal to the lower envelope of its core (i.e., for all $A \in \Sigma$, $\nu(A) = \min_{P \in \mathcal{C}(\nu)} P(A)$). Convex implies exact, which in turn implies balanced, but the converse implications are all false.

The notion of integral used for capacities is the **Choquet integral**, due to Choquet [5]. For a given Σ-measurable function $\varphi : S \to \mathbb{R}$, the Choquet integral of φ with respect to a capacity ν is defined as:

$$\int_S \varphi \, d\nu = \int_0^\infty \nu(\{s \in S : \varphi(s) \geq \alpha\}) d\alpha + \int_{-\infty}^0 [1 - \nu(\{s \in S : \varphi(s) \geq \alpha\})] d\alpha \qquad (21)$$

where the r.h.s. is a Riemann integral (which is well defined because ν is monotone). When ν is additive, (21) becomes a standard (additive) integral. In general it is seen to be monotonic, positive homogeneous and **comonotonic additive**: If $\varphi, \psi : S \to \mathbb{R}$ are non-negative and comonotonic, then $\int (\varphi + \psi) \, d\nu = \int \varphi \, d\nu + \int \psi \, d\nu$. Two functions $\varphi, \psi : S \to \mathbb{R}$ are called **comonotonic** if there are no $s, s' \in S$ such that $\varphi(s) > \varphi(s')$ and $\psi(s) < \psi(s')$.

Appendix B Cardinal Symmetry and Biseparable Preferences

In this Appendix, we prove Proposition 6. In order to make the proof as clear as possible, we first explain the notion of 'standard sequence', and then show how the latter can be used to prove the proposition.

B.1 Standard Sequences

Consider a DM whose preferences have a canonical representation V, with canonical utility index u, willingness to bet ρ, and an essential event $A \in \Sigma$. Fix a pair of consequences $v^* \succ v_*$, and consider $x^0 \in X$ such that $x^0 \succ v^*$. If there is an $x \in X$ such that

$x\,A\,v_* \succ x^0\,A\,v^*$, then by (3) and the convexity of the range of u, there is $x^1 \in X$ such that

$$x^1\,A\,v_* \sim x^0\,A\,v^*. \tag{22}$$

It is easy to verify that $x^1 \succ x^0$: If $x^0 \succcurlyeq x^1$ held, by monotonicity and biseparability, we would have $x^0\,A\,v^* \succcurlyeq x^1\,A\,v^*$ and $x^1\,A\,v^* \succ x^1\,A\,v_*$. This yields $x^0\,A\,v^* \succ x^1\,A\,v_*$, a contradiction. Assuming that there is an $x \in X$ such that $x\,A\,v_* \succ x^1\,A\,v^*$, as above we can find $x^2 \in X$ such that

$$x^2\,A\,v_* \sim x^1\,A\,v^*. \tag{23}$$

Again, $x^2 \succ x^1$. We can use the representation V to check that the equivalences in (22) and (23) translate to

$$u(x^1) - u(x^0) = \frac{1-\rho(A)}{\rho(A)}(u(v^*) - u(v_*)) = u(x^2) - u(x^1), \tag{24}$$

that is, the three points x^0, x^1, x^2, are *equidistant* in u. Proceeding in this fashion we can construct a sequence of points $\{x^0, x^1, x^2, \dots\}$ all evenly spaced in utility. Such sequence we call an **increasing standard sequence with base x^0, carrier A and mesh** (v_*, v^*). (Notice that the distance in utility between the points in the sequence is proportional to the distance in utility between v_* and v^*, which is used as the 'measuring rod'.)

Analogously, we can construct a **decreasing** standard sequence with base x^0, carrier A and mesh (v_*, v^*) where $v_* \succ x^0$. This will be a sequence starting again from x^0, but now moving in the direction of decreasing utility: For every $n \geq 0$, $v^*\,A\,x^{n+1} \sim v_*\,A\,x^n$. Henceforth, we call a **standard sequence w.r.t.** (x^0, A) any sequence $\{\bar{x}^0, \bar{x}^1, \bar{x}^2, \dots\}$ such that $\bar{x}^0 = x^0$, and there is a pair of points (above or below x^0) which provides the mesh for obtaining $\{\bar{x}^0, \bar{x}^1, \bar{x}^2, \dots\}$ as a decreasing/increasing standard sequence with carrier A.

It is simple to see how — having fixed an essential event A, and a base x^0 which is **non-extremal** in the ordering on X (i.e., there are $y, z \in X$ such that $y \succ x^0 \succ z$) — standard sequences can be used to measure the canonical utility index u of a biseparable preference (extending the scope of the method proposed by Wakker and Deneffe [29]): One just needs to construct (increasing and decreasing) standard sequences with base x^0 and finer and finer mesh. In what follows we use standard sequences and cardinal symmetry to show that equality of the u_i, $i = 1, 2$, can be verified *without* eliciting them.

B.2 Equality of Utilities: Proof of Proposition 6

The proof of Proposition 6 builds on two lemmas. The first lemma, whose simple proof we omit, shows the following: Suppose that a pair of biseparable preferences are cardinally symmetric, then for fixed non-extremal x^0 and essential events A_1 and A_2, the sets of the standard sequences (with respect to (x^0, A_1) and (x^0, A_2) respectively) of the orderings are 'nested' into each other. Stating this lemma requires some terminology and notation: Given a standard sequence $\{x^n\}$ for preference relation $\succcurlyeq_i$, we say that a sequence $\{y^m\} \subseteq X$ is a **refinement** of $\{x^n\}$ if it is itself a standard sequence, and it is such that $y^m = x^n$ whenever $m = kn$ for some $k \in \mathbb{N}$. Two canonical utility indices are subject to a **common normalization** if they take identical values on two consequences $x, y \in X$ such that $x \succ_i y$ for both i. Finally, for the rest of this section: For each $i = 1, 2$, the carrier of any standard sequence for $\succcurlyeq_i$ is a fixed essential event A_i, and $SQ(\succcurlyeq_i, x^0) \subseteq X$ denotes the set of the points belonging to some standard sequence of $\succcurlyeq_i$ with base x^0 and carrier A_i.

Lemma 27 *Suppose that $\succcurlyeq_1$, $\succcurlyeq_2$ are as assumed in Proposition 6. Fix a non-extremal $x^0 \in X$.* If *$\succcurlyeq_1$ and $\succcurlyeq_2$ are cardinally symmetric,* then *the following holds: Either every standard sequence for ordering $\succcurlyeq_1$ is a refinement of a standard sequence for $\succcurlyeq_2$, or every standard sequence for ordering $\succcurlyeq_2$ is a refinement of a standard sequence for $\succcurlyeq_1$. Hence, $SQ(\succcurlyeq_1, x^0) = SQ(\succcurlyeq_2, x^0) \equiv SQ(x^0)$.*

The second lemma shows that, because of cardinal symmetry, the result holds on $SQ(x^0)$:

Lemma 28 *Suppose that $\succcurlyeq_1$, $\succcurlyeq_2$ are as assumed in Proposition 6.* If *$\succcurlyeq_1$ and $\succcurlyeq_2$ are cardinally symmetric,* then *for any non-extremal $x^0 \in X$ and any common normalization of the two indices, $u_1(x) = u_2(x)$ for every $x \in SQ(x^0)$.*

Proof: Fix a non-extremal x^0. Suppose that x belongs to an increasing standard sequence for $\succcurlyeq_i$, $\{x^n\}$. Since the relations are cardinally symmetric, by Lemma 27 it is w.l.o.g. (taking refinements if necessary) to take the sequence to be standard for both orderings. That is, there are $v_*, v^*, w_*, w^* \in X$ such that $v^* \succ_1 v_*$, $w^* \succ_2 w_*$ and for $n \geq 0$,

$$x^{n+1} \, A_1 \, v_* \sim_1 x^n \, A_1 \, v^*,$$

and analogously for $\succcurlyeq_2$ (with w replacing v). Moreover, there is $n \geq 0$ such that $x = x^n$. Choose x^m for some $m > n$, and take positive affine transformations of the two canonical utility functions so as to obtain $u_1(x^0) = u_2(x^0) = 0$ and $u_1(x^m) = u_2(x^m) = 1$. All points in the sequence are evenly spaced for both preferences (cf. Eq. (24)). Hence we have $u_1(x^n) = u_2(x^n) = n/m$. The case in which x belongs to a decreasing standard sequence is treated symmetrically. Finally, we have the immediate observation that if $u_1(x) = u_2(x)$ for one common normalization, the equality holds for every common normalization. ∎

Proof of Proposition 6: The 'if' part follows immediately from the canonical representation. We now prove the 'only if.' Start by fixing a non-extremal x^0 and adding a constant to both indices, so that $u_1(x^0) = u_2(x^0) = 0$. Suppose that (after this transformation) there is $x \in X$ such that $u_1(x) \neq u_2(x)$. By relabelling if necessary, assume that

$u_1(x) = \alpha > \beta = u_2(x)$. There are different cases to consider, depending on where α and β are located.

Suppose first that $\beta \geq 0$. Choose $v^* \in X$ such that $x^0 \succ_1 v^*$ and further transform the utilities so that $\bar{u}_1(v^*) = \bar{u}_2(v^*) = -1$, to obtain $\bar{u}_1(x) = \bar{\alpha} > \bar{\beta} = \bar{u}_2(x)$. Choose $\varepsilon > 0$ such that $\bar{\alpha} - \bar{\beta} > \varepsilon$. By the connectedness of the range of each u_i and Lemma 27, there are $v_*, w_* \in X$ such that (v_*, v^*) and (w_*, v^*) generate the same standard sequence $\{x^n\}$ and

$$\bar{u}_1(x^{n+1}) - \bar{u}_1(x^n) = \bar{u}_2(x^{n+1}) - \bar{u}_2(x^n) < \varepsilon.$$

So the 'length' of the utility interval between each element in the increasing standard sequence is smaller than the distance between $\bar{\alpha}$ and $\bar{\beta}$. We also proved in Lemma 28 that for each element in the standard sequence, we have equality of the utilities (since we imposed a common normalization). Hence there must be $n \geq 0$ such that $\bar{u}_1(x^n) = \bar{u}_2(x^n) = \gamma \in (\bar{\beta}, \bar{\alpha})$. We then have

$$\bar{u}_1(x^n) > \bar{u}_1(x) \Leftrightarrow x^n \succ_1 x \quad \text{and} \quad \bar{u}_2(x^n) < \bar{u}_2(x) \Leftrightarrow x^n \prec_2 x,$$

which contradicts the assumption of ordinal equivalence.

The case in which $\alpha \leq 0$ is treated symmetrically. If, finally, $\alpha > 0 > \beta$ then, using an argument similar to the one just presented, one can find $\bar{x} \in X$ such that $u_1(\bar{x}) = u_2(\bar{x}) \in (0, \alpha)$ and obtain a similar contradiction. This shows that $u_1(x) = u_2(x)$ for every $x \in X$. ∎

Appendix C Proofs for Sections 3 to 5

C.1 Section 3

Proof of Theorem 10: We first state without proof an immediate result:

Lemma 29 *Two preference relations $\succcurlyeq_1$ and $\succcurlyeq_2$ satisfying Eqs. (7) and (8) are ordinally equivalent.*

Given this lemma, if $\succcurlyeq_1$ and $\succcurlyeq_2$ have essential events the result follows immediately from Proposition 6. If, say, relation $\succcurlyeq_i$ does not have essential events, any ordinal transformations of u_i is still a canonical utility. Since the two preferences are ordinally equivalent by the lemma, it is then w.l.o.g. to use u_j ($j \neq i$) to represent both of them. ∎

Proof of Theorem 12: We first prove that $\mathcal{D}(\succcurlyeq) \subseteq \mathcal{M}(\succcurlyeq)$. Given a canonical representation V of $\succcurlyeq$ with canonical utility u, suppose that $P \in \mathcal{D}(\succcurlyeq)$, and consider the relation

$\succcurlyeq$ induced by P and u. We want to show that $\succcurlyeq$ is more ambiguity averse than $\succcurlyeq$. Since $P \in \mathcal{D}(\succcurlyeq)$, $\int u(f)\, dP \geq V(f)$ for all $f \in \mathcal{F}$, so that for every $x \in X$ and $f \in \mathcal{F}$,

$$u(x) \geq \int_S u(f(s))\, P(ds) \Longrightarrow V(x) \geq V(f),$$

where the implication follows from the definition $u(x) = V(x)$ for all $x \in X$. This proves that (7) holds. Similarly one shows the validity of (8). Part *(B)* of Definition 7 is immediate: If $\succcurlyeq$ and $\succcurlyeq$ have essential events, then the result follows from Proposition 6. Hence $\succcurlyeq \in \mathcal{R}(\succcurlyeq)$, or in other words $P \in \mathcal{M}(\succcurlyeq)$.

We now prove the opposite inclusion $\mathcal{D}(\succcurlyeq) \supseteq \mathcal{M}(\succcurlyeq)$. Suppose that $P \in \mathcal{M}(\succcurlyeq)$. Let $\succcurlyeq$ be the benchmark preference corresponding to P, and let u' be the canonical utility index of $\succcurlyeq$. Since $\succcurlyeq$ is a benchmark for $\succcurlyeq$, we have for every $x \in X$ and $f \in \mathcal{F}$,

$$u'(x) \geq \int_S u'(f(s))\, P(ds) \Longrightarrow u(x) \geq V(f), \tag{25}$$

and the same with strict inequality. We have to show that $P \in \mathcal{D}(\succcurlyeq)$. By Theorem 10, it is w.l.o.g. to take $u = u'$. Hence, (25) implies that $\int u(f)\, dP \geq V(f)$ for all $f \in \mathcal{F}$, and so $P \in \mathcal{D}(\succcurlyeq)$. ∎

Proof of Corollary 13: By Theorem 12, $\mathcal{M}(\succcurlyeq) = \mathcal{D}(\succcurlyeq)$. Let $P \in \mathcal{D}(\succcurlyeq)$. For every $A \in \Sigma$ and $x^* \succ x_*$, consider the act $f = x^*\, A\, x_*$. Normalizing $u(x^*) = 1$ and $u(x_*) = 0$, we have

$$P(A) = \int_S u(f(s))\, P(ds) \geq \int_S u(f(s))\, \nu(ds) = \nu(A),$$

and so $P \in \mathcal{C}(\nu)$. This implies $\mathcal{D}(\succcurlyeq) \subseteq \mathcal{C}(\nu)$. The converse inclusion is trivial, since $P \in \mathcal{C}(\nu)$ implies $\int u(f)\, dP \geq \int u(f)\, d\nu$ for all $f \in \mathcal{F}$. ∎

Proof of Corollary 14: We are done if we show that for all $f, g \in \mathcal{F}$,

$$f \succcurlyeq g \Longleftrightarrow \min_{P \in \mathcal{D}(\succcurlyeq)} \int_S u(f(s))\, P(ds) \geq \min_{P \in \mathcal{D}(\succcurlyeq)} \int_S u(g(s))\, P(ds). \tag{26}$$

This follows from the fact that there exists a *unique* weak*-compact and convex set C representing $\succcurlyeq$. $\mathcal{D}(\succcurlyeq)$ is clearly weak*-compact (so that the minimum in (26) is well defined) and convex. Hence, if (26) holds $C = \mathcal{D}(\succcurlyeq)$, and by Theorem 12, $\mathcal{D}(\succcurlyeq) = \mathcal{M}(\succcurlyeq)$.

To prove (26), suppose there are $f, g \in \mathcal{F}$ such that

$$\min_{P \in C} \int u(f)\, dP \geq \min_{P \in C} \int u(g)\, dP \quad \text{and} \quad \min_{P \in \mathcal{D}(\succcurlyeq)} \int u(f)\, dP < \min_{P \in \mathcal{D}(\succcurlyeq)} \int u(g)\, dP.$$

Let $P^* \in \arg\min\{\int_S u(f(s))\, P(ds) : P \in \mathcal{D}(\succcurlyeq)\}$. Since $C \subseteq \mathcal{D}(\succcurlyeq)$, we have:

$$\min_{P \in C} \int_S u(f(s))\, P(ds) \leq \int_S u(f(s))\, P^*(ds) < \min_{P \in \mathcal{D}(\succcurlyeq)} \int_S u(g(s))\, P(ds) \leq \min_{P \in C} \int_S u(g(s))\, P(ds),$$

a contradiction. Similarly, one shows that there cannot be $f, g \in \mathcal{F}$ such that the preference based on $\mathcal{D}(\succcurlyeq)$ prefers weakly f to g, while $g \succ f$. This shows that Eq. (26) holds, concluding the proof. ∎

Proof of Proposition 15: That every SEU preference is ambiguity neutral follows immediately from two applications of Theorem 13. As for the converse: If $\succcurlyeq$ is both ambiguity averse and ambiguity loving, there are a SEU preference relation $\geqslant_1$ (represented by probability P_1) such that $\succcurlyeq$ is more ambiguity averse than $\geqslant_1$, and a SEU preference relation $\geqslant_2$ (represented by probability P_2) which is more ambiguity averse than $\succcurlyeq$. Applying Definition 7 twice, we obtain that for every $f \in \mathcal{F}$ and $x \in X$,

$$x \geqslant_1 f \Rightarrow x \geqslant_2 f \qquad \text{and} \qquad x >_1 f \Rightarrow x >_2 f.$$

We show that $\geqslant_1$ and $\geqslant_2$ are cardinally symmetric. This requires first showing that if $\geqslant_2$ has an essential event, so must $\succcurlyeq$. Suppose that $A \in \Sigma$ is essential for $\geqslant_2$, so that for some $x \succ y$ (remember that $\succcurlyeq$ and $\geqslant_1$ and $\geqslant_2$ are all ordinally equivalent), $x >_2 x\,A\,y >_2 y$. Using the contrapositive of (7), we then have $x\,A\,y \succ y$. Since $\geqslant_2$ is a SEU preference, A^c is also $\geqslant_2$-essential, similarly implying $x\,A^c\,y \succ y$. Now, suppose that $\succcurlyeq$ has no essential event. Because of the preferences we just derived, we must have both $x \sim x\,A\,y$ and $x \sim x\,A^c\,y$. This is impossible since $\geqslant_1 \in \mathcal{R}(\succcurlyeq)$, for the contrapositive of (8) then yields $x\,A\,y \geqslant_1 x$, which implies $P_1(A) = 1$, and $x\,A^c\,y \geqslant_1 x$, which implies $P_1(A) = 0$. This gives us a contradiction, so that $\succcurlyeq$ must have an essential event if $\geqslant_2$ does. Hence, $\geqslant_2$ and $\succcurlyeq$ have essential events, and they are cardinally symmetric by assumption. Similarly one shows that $\geqslant_1$ and $\succcurlyeq$ have essential events and are cardinally symmetric. It is now immediate to check that these facts imply that $\geqslant_1$ and $\geqslant_2$ are cardinally symmetric. We thus conclude that $\geqslant_2$ is more ambiguity averse than $\geqslant_1$. Mimicking the last part of the proof of Theorem 12, we then show that then $P_1 \geq P_2$, which immediately implies $P_1 = P_2$, so that $\geqslant_1 = \geqslant_2 \equiv \geqslant$. Thus $\succcurlyeq$ is both more and less ambiguity averse than $\geqslant$, which immediately implies $\succcurlyeq = \geqslant$. ∎

Proof of Theorem 17: Part (i) follows immediately along the lines of the proofs of Theorem 12 and Corollary 13. As for part (ii), it is similarly immediate to show that if $\succcurlyeq_2$ is more ambiguity averse than $\succcurlyeq_1$, then $C_1 \subseteq \mathcal{D}(\succcurlyeq_2)$ and $u_1 \approx u_2$. We show the converse. Let V_1 and V_2 denote the canonical representations of $\succcurlyeq_1$ and $\succcurlyeq_2$, and w.l.o.g. assume that $u_1 = u_2 = u$. Then $C_1 \subseteq \mathcal{D}(\succcurlyeq_2)$ implies that for every $f \in \mathcal{F}$ and every $P \in C_1$, $V_2(f) \leq \int u(f)\, dP$. Hence, using the fact that $\succcurlyeq_1$ is MEU, we find

$$V_2(f) \leq \min_{P \in C_1} \int_S u(f(s))\, P(ds) = V_1(f),$$

which immediately yields the desired result. ∎

C.2 Section 4

Proof of Proposition 22: Let $\geqslant \in \mathcal{R}(\succcurlyeq)$ and set $\Theta \equiv \{A \in \Sigma : \rho(A) + \rho(A^c) = 1\}$. If $A \in \Lambda_{\geqslant}$ for all $x \in X$ we have

$$u(x) = P(A) \iff u(x) = \rho(A)$$
$$u(x) = P(A^c) \iff u(x) = \rho(A^c),$$

and so $\rho(A) = P(A)$ and $\rho(A^c) = P(A^c)$. This implies that $A \in \Theta$, so that $\Lambda_{\geqslant} \subseteq \Theta$.

Now, if $A \in \Theta$ we have

$$\rho(A) = P(A) \text{ and } \rho(A^c) = P(A^c). \tag{27}$$

In order to show that $A \in \Lambda_{\geqslant}$, we need to show that any act measurable w.r.t. the partition $\{A, A^c\}$ is in $\mathcal{H}_{\geqslant}$. This follows from (27), as for every $x, y \in X$ we have $V(x\,A\,y) = V_{\geqslant}(x\,A\,y)$. Thus $\Theta \subseteq \Lambda_{\geqslant}$, which concludes the proof. ∎

C.3 Section 5

Proof of Proposition 23: Suppose, to the contrary, that ν agrees with (17). If Eq. (15) holds then $P(R) = \nu(R)$ and $P(B, Y) = \nu(B, Y)$ for all $P \in \mathcal{C}(\nu)$, so that we have

$$P\,(B, Y) = \nu\,(B, Y) < \nu\,(B, R) \leq P(B, R).$$

In turn, this implies $P(Y) < P(R)$, yielding $\nu(Y) \leq P(Y) < P(R) = \nu(R)$. Hence $\nu(Y) < \nu(R)$, contradicting (17). ∎

Proof of Proposition 24: Every ν which satisfies (18) is such that $\mathcal{C}(\nu) \neq \emptyset$. For, the measure P such that $P(R) = P(B) = P(Y) = 1/3$ belongs to $\mathcal{C}(\nu)$. This proves that all preferences satisfying (18) are ambiguity averse.

As to the converse, let $\succcurlyeq$ be ambiguity averse, i.e. $\mathcal{C}(\nu) \neq \emptyset$. Let $P \in \mathcal{C}(\nu)$. Assume first that $\nu(B) = \nu(Y) > \nu(R)$. Since $P(B) \geq \nu(B)$ and $P(Y) \geq \nu(Y)$,

$$P(B) + P(R) + P(Y) > \nu(B) + \nu(R) + \nu(Y) > 1,$$

a contradiction. Assume now $\nu(B, Y) < \nu(B, R) = \nu(R, Y)$. This implies $P(B, Y) < P(B, R)$ and $P(B, Y) < P(R, Y)$, so that $P(Y) < P(R)$, $P(B) < P(R)$, and $P(B) + P(R) + P(Y) < 1$, a contradiction. ∎

References

[1] Frank J. Anscombe and Robert J. Aumann. A definition of subjective probability. *Annals of Mathematical Statistics*, 34:199–205, 1963.

[2] Kenneth J. Arrow. The theory of risk aversion. In *Essays in the Theory of Risk-Bearing*, chapter 3. North-Holland, Amsterdam, 1974. (Part of the Yriö Jahnssonin Säätio lectures in Helsinki, 1965).

[3] Ramon Casadesus-Masanell, Peter Klibanoff, and Emre Ozdenoren. Maxmin expected utility over Savage acts with a set of priors. *Journal of Economic Theory*, 92:33–65, 2000.

[4] Alain Chateauneuf and Jean-Marc Tallon. Diversification, convex preferences and non-empty core. Mimeo, Université Paris I, July 1998.

[5] Gustave Choquet. Theory of capacities. *Annales de l'Institut Fourier (Grenoble)*, 5:131–295, 1953.

[6] Bruno de Finetti. Sulla preferibilità. *Giornale degli Economisti e Annali di Economia*, 6:3–27, 1952.

[7] Daniel Ellsberg. Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, 75:643–669, 1961.

[8] Larry G. Epstein. A definition of uncertainty aversion. *Review of Economic Studies*, 66:579–608, 1999.

[9] Larry G. Epstein and Tan Wang. Intertemporal asset pricing under Knightian uncertainty. *Econometrica*, 62:283–322, 1994.

[10] Larry G. Epstein and Jiankang Zhang. Subjective probabilities on subjectively unambiguous events. *Econometrica*, forthcoming.

[11] Peter C. Fishburn. The axioms and algebra of ambiguity. *Theory and Decision*, 34:119–137, 1993.

[12] Paolo Ghirardato and Jonathan N. Katz. Indecision theory: Explaining selective abstention in multiple elections. Social Science Working Paper 1106, Caltech, November 2000. `http://masada.hss.caltech.edu/~paolo/voting.pdf`.

[13] Paolo Ghirardato and Massimo Marinacci. Risk, ambiguity, and the separation of utility and beliefs. Social Science Working Paper 1085, Caltech, March 2000. (Revised: August 2000); `http://masada.hss.caltech.edu/~paolo/canonica4wp.pdf`.

[14] Paolo Ghirardato and Massimo Marinacci. A subjective definition of ambiguity. Work in progress, Caltech and Università di Torino, 2000.

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

[16] Lars P. Hansen, Thomas Sargent, and Thomas D. Tallarini. Robust permanent income and pricing. *Review of Economic Studies*, 66:873–907, 1999.

[17] Yakar Kannai. The core and balancedness. In Robert J. Aumann and Sergiu Hart, editors, *Handbook of Game Theory*, pages 355–395. North–Holland, Amsterdam, 1992.

[18] David Kelsey and Shasikanta Nandeibam. On the measurement of uncertainty aversion. Mimeo, University of Birmingham, September 1996.

[19] David H. Krantz, R. Duncan Luce, Patrick Suppes, and Amos Tversky. *Foundations of Measurement: Additive and Polynomial Representations*, volume 1. Academic Press, San Diego, 1971.

[20] Mark J. Machina and David Schmeidler. A more robust definition of subjective probability. *Econometrica*, 60:745–780, 1992.

[21] Aldo Montesano and Francesco Giovannoni. Uncertainty aversion and aversion to increasing uncertainty. *Theory and Decision*, 41:133–148, 1996.

[22] Sujoy Mukerji. Ambiguity aversion and incompleteness of contractual form. *American Economic Review*, 88:1207–1231, 1998.

[23] Klaus Nehring. Capacities and probabilistic beliefs: A precarious coexistence. *Mathemathical Social Sciences*, 38:197–213, 1999.

[24] John W. Pratt. Risk aversion in the small and in the large. *Econometrica*, 32:122–136, 1964.

[25] Leonard J. Savage. *The Foundations of Statistics*. Wiley, New York, 1954.

[26] David Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57:571–587, 1989.

[27] Amos Tversky and Peter P. Wakker. Risk attitudes and decision weights. *Econometrica*, 63:1255–1280, 1995.

[28] Peter P. Wakker. *Additive Representations of Preferences*. Kluwer, Dordrecht, 1989.

[29] Peter P. Wakker and Daniel Deneffe. Eliciting von Neumann-Morgenstern utilitites when probabilities are distorted or unknown. *Management Science*, 42:1131–1150, 1996.

[30] Menachem E. Yaari. Some remarks on measures of risk aversion and on their uses. *Journal of Economic Theory*, 1:315–329, 1969.

[31] Jiankang Zhang. Subjective ambiguity, probability and capacity. Mimeo, University of Toronto, October 1996.