Available online at www.sciencedirect.com
ScienceDirect
Journal of Economic Theory 139 (2008) 1–24
JOURNAL OF Economic Theory
www.elsevier.com/locate/jet

Small worlds: Modeling attitudes toward sources of uncertainty

Soo Hong Chew[a], Jacob S. Sagi[b,*]

[a] *Department of Economics, Hong Kong University of Science and Technology, Hong Kong*
[b] *Owen Graduate School of Business, Vanderbilt University, USA*

Received 26 June 2006; final version received 2 July 2007
Available online 20 September 2007

Abstract

We introduce the concept of a *conditional small world event domain*—an extension of Savage's [The Foundations of Statistics, Wiley, New York, 1954] notion of a 'small world'—as a self-contained collection of comparable events. Under weak behavioral conditions we demonstrate probabilistic sophistication in any small world event domain without relying on monotonicity or continuity. Probabilistic sophistication within, though not necessarily across, small worlds provides a foundation for modeling a decision maker that has source-dependent risk attitudes. This also helps formalize the idea of source preference and suggests an interpretation of ambiguity aversion, often associated with Ellsberg-type behavior, in terms of comparative risk aversion across small worlds.

JEL classification: D11; D81

Keywords: Uncertainty; Risk; Ambiguity; Decision theory; Non-expected utility; Utility representation; Probabilistic sophistication; Ellsberg paradox

1. Introduction

In what has come to be known as Ellsberg's [6] two-urn problem, a related version of which was proposed in Keynes [20], one urn contains 50 red and 50 black balls while the second contains an unspecified combination of the two. It is observed that individuals are indifferent to betting on a red ball versus a black ball from the same urn, and yet prefer to bet on a ball (of either color) from the urn with the known mixture. Such betting preferences are not consistent with

* Corresponding author.
E-mail addresses: chew.sh@ust.hk (S.H. Chew), Jacob.Sagi@Owen.Vanderbilt.edu (J.S. Sagi).

0022-0531/$ - see front matter © 2007 Elsevier Inc. All rights reserved.
doi:10.1016/j.jet.2007.07.004

probabilistic sophistication: A decision maker is said to be probabilistically sophisticated if her choice behavior reflects probabilistic beliefs (in the sense that events are distinguished only by subjective probabilities that are uniquely assigned to them). In the two-urn experiment, decision makers appear to prefer betting on a 50% event from one urn rather than a 50% event from another urn. This suggests that each urn can be associated with a distinct restricted domain of events—a sufficiently simple source of uncertainty that we refer to as a 'small world'—such that the decision maker is probabilistically sophisticated within each small world, but equally likely complementary events in one small world may not be comparable with equally likely complementary events in another small world. [1]

The observed choice behavior over Ellsbergian urns has inspired a substantial literature in axiomatic models of decision making. [2] Similar non-comparability of equally likely events from distinct sources of uncertainty, not specific to situations where information about objective probabilities is imprecise, appears pervasive. For instance, in playing Lotto, customers may prefer selecting their own numbers rather than having them picked by a computer. Betting on the outcome of tossing one's own coin may be preferable to doing so with a coin from someone unfamiliar. Having one's favorite celebrity determine the color mixture in the unknown urn may render betting on it more attractive. There is also a growing literature recognizing the presence of source-dependent preferences in decision making. [3]

The examples mentioned provide vivid instances of deviation from global probabilistic sophistication as well as the seeming presence of restricted collections of events over which the decision maker may exhibit probabilistic sophistication. This leads naturally to the following question:

> What are appropriate subjective conditions that characterize self-contained *small world event domains*, such that the decision maker's preference over acts restricted to any one domain exhibits probabilistic sophistication?

The Machina and Schmeidler [25] axioms can be used to define collections of 'subjectively unambiguous events,' as is done in Epstein and Zhang [8] and Kopylov [23]. The approaches used by these authors are, however, only concerned with determining a single source of uncertainty consistent with probabilistic sophistication (namely, the collection of subjectively unambiguous events). In addition, one could question whether assumptions such as monotonicity, continuity, or comparative likelihood, used in both Savage and Machina and Schmeidler are pertinent to 'small worlds' probabilistic sophistication, per se.

In this paper, we address the above question using a different approach based on a recent paper by Chew and Sagi ([4], CS06 henceforth). The latter derives a characterization of probabilistic sophistication that eschews the standard Savage axioms of monotonicity (P3), continuity (P6), comparative likelihood (P4), and the Sure Thing Principle (P2). [4] The approach of CS06 states choice-behavioral axioms in terms of the exchangeability of events. Two events are said to be

[1] Savage [33] recognized that decision making often takes place within 'small worlds' but, in contrast with us, adopted the view that events can only be distinguished based on their likelihood and not on their source.

[2] Some key references include Schmeidler [34], Gilboa and Schmeidler [14], and Nakamura [26]. A number of works posit the Knightian [22] distinction between risk and uncertainty and classify events as being either subjectively ambiguous or subjectively unambiguous (see [8,12,21,13,36,28,29]).

[3] See, for example, Heath and Tversky [17], Fox and Tversky [11], Keppe and Weber [19], Tversky and Wakker [38], Schroder and Skiadas [35], Skiadas [37], as well as an early reference in Fellner [10]. Hsu et al. [18] link source preference to brain activation in the orbitofrontal cortex using functional brain imaging.

[4] Machina and Schmeidler [25] and Grant [15] give up the Sure Thing Principle and retain limited or strengthened versions of the remaining axioms.

exchangeable if the decision maker is always indifferent to permuting their payoffs. [5] This leads to the definition of an event E as being 'at least as likely as' another event E' if it contains a subevent that is exchangeable with E'. Completeness of this relation is a key assumption, Axiom C, in CS06. When the state space is finite, Axiom C along with two other conditions, guarantees that the decision maker is able to partition the state space into equally likely events, and that she is indifferent to betting on events from two different equi-partitions as long as the two partitions have the same number of elements. In investigating departures from global probabilistic sophistication one is naturally led to relaxing Axiom C. By contrast, it is not clear which of the Machina and Schmeidler [25] axioms one ought to relax to depart from probabilistic sophistication. CS06, for instance, note that Machina and Schmeidler's [25] comparative likelihood axiom (Axiom P4*)—the axiom that could be deemed the natural candidate for removal in order to relax probabilistic sophistication—is incompatible with a large class of explicitly probabilistically sophisticated preferences (thus relaxing P4* does not immediately imply the absence of probabilistic sophistication).

In Section 3 we dispense with Axiom C and define a *conditional small world event domain* (or conditional small world for short) as a collection, $\mathcal{A}$, of events that is self-contained in the following key senses: given any pair of events in the collection $\mathcal{A}$, one event contains a subevent that is exchangeable with the other; the union of all collection elements is in the collection; if E is the 'more likely' of two events in $\mathcal{A}$, then both the subevent of E that is exchangeable with E', as well as the remainder of E, are in $\mathcal{A}$; finally, if $\{E_i\}_{i=1}^n$ and $\{E_i'\}_{i=1}^n$ are partitions in $\mathcal{A}$ such that E_i is exchangeable with E_i' for all $i = 1, \ldots, n$, then the decision maker is indifferent between an act that pays x_i on E_i, $i = 1, \ldots, n$, and the act that instead pays x_i on E_i', $i = 1, \ldots, n$. Each conditional small world models a single, subjectively viewed and self contained homogeneous source of uncertainty. This description of a conditional small world renders it a λ-system whose universal set may be a proper subset of the full state space (hence the term 'conditional'). [6] If the universal set of $\mathcal{A}$ is the full state space, then our notion is analogous to Savage's concept of a small world (see [33, pp. 9, 27, and 81]). In such a situation, we say that $\mathcal{A}$ is an *unconditional small world*. Finally, we say that a small world domain is *maximal* if the collection resulting from adding any other collection of non-null events to the domain is not compatible with (i)–(iv). Small world event domains can be viewed as subjectively distinct sources of uncertainty.

We impose two axioms: a set of non-null, pairwise disjoint and exchangeable events must be finite, and any event that is inconsequential, when it is 'added' to either of the two events in a conditional binary bet, is null. Our main theorem establishes that these two conditions suffice to deliver small world or source-specific probabilistic sophistication. Proof of this is non-trivial, and made especially so because of the non-algebraic nature of the event collection making up a conditional small world. The fact that we do not require monotonicity or continuity can be viewed as extending the contribution of Epstein and Zhang [8], whose construction for probabilistic sophistication over a λ-system is based on the Machina and Schmeidler [25] axioms, and is specifically intended for modeling a single source of uncertainty. By contrast, our construction is readily applicable to probabilistic sophistication within individual sources of uncertainty, thereby accommodating the possibility of the decision maker having distinct risk attitudes for lotteries defined on different small worlds. This suggests an interpretation of ambiguity aversion in terms of comparative risk aversion across small worlds (within the same individual).

[5] The term exchangeability is also used, albeit to refer to different concepts, by de Finetti [5, see Chapters 3 and 5] and Ramsey [31].

[6] A λ-system, defined in Section 3, is a collection of events suitable for a natural definition of a probability measure.

Section 2 provides basic definitions of our primitives, Section 3 states our axioms and results, investigates some examples of small world event domains, and considers possible utility representations. Section 4 discusses the relationship between our work and that of Epstein and Zhang [8].

2. Preliminaries

2.1. Exchangeability and comparability

Let Ω be a space whose elements correspond to all states of the world. Let X be a set of payoffs and Σ an algebra on Ω. Elements of Σ are events. If $e, E \in \Sigma$ and $e \subseteq E$, then we say that e is a *subevent* of E. The set of simple acts, $\mathcal{F}$, comprises all Σ-adapted and X-valued functions over Ω that have a finite range. As is customary, $x \in X$ is identified with the *constant act* that pays x in every state. Throughout the paper we assume that the decision maker has a non-degenerate binary preference relation, $\succeq$, on $\mathcal{F}$ (Savage's P1 and P5).[7]

For any collection of pairwise disjoint events, $E_1, E_2, \ldots, E_n \subset \Omega$, and $f_1, f_2, \ldots, f_n, g \in \mathcal{F}$, let $f_1 E_1 f_2 E_2 \ldots f_n E_n g$ denote the act that pays $f_i(\omega)$ if the true state, $\omega \in \Omega$, is in E_i, and pays $g(\omega)$ otherwise. We say that $E \in \Sigma$ is *null* if $fEh \sim gEh\ \forall f, g, h \in \mathcal{F}$.

The following two definitions are from CS06:

Definition 1 (*Event exchangeability*). For any pair of disjoint events $E, E' \in \Sigma$, $E \approx E'$ if for any $x, x' \in X$ and $f \in \mathcal{F}$, $xEx'E'f \sim x'ExE'f$.

Whenever $E \approx E'$ we will say that E and E' are *exchangeable*. Note that all null events are exchangeable. Exchangeability can be viewed as expressing a notion of 'equal likelihood': two events are 'equally likely' if the decision maker is indifferent to a permutation of their payoffs. As observed in CS06, $\approx$ is not necessarily transitive, and therefore not an equivalence relation.

Intuitively, an event is 'at-least-as-likely' as any of its subevents. Exchangeability motivates a similar comparison across disjoint events, $E, E' \in \Sigma$: if a subevent of E is exchangeable with E', then it is also natural to view E as 'at-least-as-likely' as E'. Building on this, we have the following exchangeability based relation between any two events.

Definition 2 (*Exchangeability-based comparative likelihood*). For any events, $E, E' \in \Sigma$, $E \succeq^C E'$ whenever $E \backslash E'$ contains a subevent, e, that is exchangeable with $E' \backslash E$. Moreover, e is referred to as a *comparison event*.

For any $E, E' \in \Sigma$, we say that E and E' are *comparable* whenever $E \succsim^C E'$ or $E' \succsim^C E$. Finally, define $E \succ^C E'$ whenever $E \succsim^C E'$ and it is not the case that $E' \succsim^C E$. Likewise, define $\sim^C$ as the symmetric part of $\succsim^C$.

Definition 3. A λ-system in Σ is a collection of events, $\mathcal{A} \subseteq \Sigma$ that satisfies the following properties:

(i) The event $\widehat{\mathcal{A}} \equiv \bigcup_{E \in \mathcal{A}} E$ is in $\mathcal{A}$.
(ii) If $E \in \mathcal{A}$ then $\widehat{\mathcal{A}} \backslash E \in \mathcal{A}$.
(iii) If $E, E' \in \mathcal{A}$ are disjoint, then $E \cup E' \in \mathcal{A}$.

---

[7] As usual, $\succ$ (resp. $\sim$) is the asymmetric (resp. symmetric) part of $\succeq$. Under Savage's P1, $\succeq$ is a weak order on $\mathcal{F}$, while P5 asserts that there exists $f, g \in \mathcal{F}$ such that $f \succ g$.

We call $\widehat{\mathcal{A}} \equiv \bigcup_{E \in \mathcal{A}} E$ the *envelope* of $\mathcal{A}$ and emphasize that, contrary to the standard definition of a λ-system, $\widehat{\mathcal{A}}$ need not coincide with Ω. Our definition is also different from the usual one in that we do not require the countable union of increasing (in the sense of set inclusion) events in $\mathcal{A}$ to be in $\mathcal{A}$.

Definition 4. A finitely additive probability measure, μ, defined on a λ-system, $\mathcal{A} \subseteq \Sigma$, represents $\succcurlyeq^C$ on $\mathcal{A}$ if and only if for every $A, B \in \mathcal{A}$, $A \succcurlyeq^C (\succ^C) B \Leftrightarrow \mu(A) \geqslant (>) \mu(B)$.

We say that $f \in \mathcal{F}$ is *adapted* to $\mathcal{A} \subseteq \Sigma$, whenever $f^{-1}(x) \cap \widehat{\mathcal{A}} \in \mathcal{A}$ for every $x \in X$. We say that f and f' induce *exchangeable partitions* in $\mathcal{A}$ whenever both are adapted to $\mathcal{A}$ and $f^{-1}(x) \cap \widehat{\mathcal{A}} \sim^C f'^{-1}(x) \cap \widehat{\mathcal{A}}$ for each $x \in X$.

The set of probability distributions, or 'lotteries', with payoffs in X is denoted Δ. For any probability measure μ on a λ-system $\mathcal{A} \subseteq \Sigma$, and act $f \in \mathcal{F}$, the lottery induced by the act, $f \in \mathcal{F}$ with respect to μ on $\mathcal{A}$ $L_{\mu,f} \equiv \{(\mu(f^{-1}(x) \cap \widehat{\mathcal{A}}), x) | x \in X\}$. The set of such lotteries is denoted $\Delta(\mu) \subseteq \Delta$. The lottery $\alpha L \oplus (1-\alpha)L'$ is defined to be an independent probabilistic α-mixture of the lotteries L and L'. We abuse notation by identifying a distribution with unit mass at $x \in X$ with x.

An *atom* is an event that cannot be partitioned into two or more non-null subevents. We say that μ is *purely and uniformly atomic* whenever the union of all atoms has unit measure and all atoms have equal measure. μ is *convex-ranged* if for every $\alpha \in [0, 1]$ and $A \in \mathcal{A}$ there is a subevent of A in $\mathcal{A}$ with $\mu(a) = \alpha\mu(A)$. Finally, we say that μ is *solvable* if for every $A, B \in \mathcal{A}$, $\mu(A) \geqslant \mu(B)$ implies the existence of a subevent of A in $\mathcal{A}$ with $\mu(a) = \mu(B)$. Requiring μ to be solvable is weaker than requiring it to be convex-ranged.

3. Sources of uncertainty, homogeneous collections of events, and small world event domains

Savage [33, pp. 9, 27, 81] discusses how decision making tends to take place in a 'small world'—events relevant to a particular decision situation, that partition the state space. While, according to Savage, one can always make a likelihood comparison between events in distinct small worlds, we explicitly incorporate the possibility that this is not the case and allow for endogenously defined distinct small worlds or sources of uncertainty. Moreover, we also consider conditional small worlds in which the events potentially partition a set strictly smaller than the full space. [8] We use the terms '(conditional) small world domain', '(conditional) small world', and 'source of uncertainty' interchangeably. [9]

In identifying collections of events suitable for a restricted notion of probabilistic sophistication, we are led by the following intuition: first, as in the case of global probabilistic sophistication, every event in the collection should be comparable with every other event; second, likelihood is generally defined relative to some benchmark event—in the case of global probabilistic sophistication

[8] Besides being useful in applications, the notion of *conditional* small worlds anticipates an inter-temporal setting in which belief updating can be studied.

[9] Gyntelberg and Hansen [16] also refer to exogenously defined small world domains of events, and derive expected utility preferences within these.

the benchmark event is Ω. Thus the collection, say $\mathcal{A}$, should contain a 'universal' event which we take to be its envelope, $\widehat{\mathcal{A}}$. Consider two events, E and E', in the collection, $\mathcal{A}$, that can potentially be described via a probability measure relative to $\widehat{\mathcal{A}}$. If $E \succeq^C E'$ then $E\backslash E'$ contains a subevent, say e, that is 'as likely as' $E'\backslash E$. Thus if the likelihood (relative to $\widehat{\mathcal{A}}$) of $E' = (E \cap E') \cup (E'\backslash E)$ is known, it should be equal to that of $\xi \equiv (E \cap E') \cup e$, and it seems sensible to require $\xi \in \mathcal{A}$. If, in addition, the likelihood (relative to $\widehat{\mathcal{A}}$) of E is known, then one can readily calculate the likelihood of $E\backslash\xi = E\backslash(E' \cup e)$, which should therefore also be in $\mathcal{A}$. Finally, the decision maker ought to be indifferent between two acts that induce exchangeable partitions of $\mathcal{A}$ and award the same payoffs outside of $\mathcal{A}$. Given these considerations, we characterize collections over which the decision maker may be probabilistically sophisticated as follows:

Definition 5. A collection of events, $\mathcal{A} \subseteq \Sigma$ is **homogeneous** if it satisfies the following:

(i) If $E, E' \in \mathcal{A}$, then E and E' are comparable.
(ii) $\widehat{\mathcal{A}} \in \mathcal{A}$ and is non-null.
(iii) For any $E, E' \in \mathcal{A}$ such that $E \succeq^C E'$, there is a comparison event, $e \subseteq E\backslash E'$, such that $(E \cap E') \cup e \in \mathcal{A}$.
(iv) If $e, E \in \mathcal{A}$ and $e \subseteq E$ then $E\backslash e \in \mathcal{A}$.
(v) If f and f' induce exchangeable partitions on $\mathcal{A}$, then for any $g \in \mathcal{F}$, $f\, \widehat{\mathcal{A}}\, g \sim f'\, \widehat{\mathcal{A}}\, g$.

If every event is comparable to every other event, then Σ itself is homogeneous and probabilistic sophistication follows from the axioms imposed in CS06. Part (iv) of the definition implies that the complement of $E \in \mathcal{A}$ is in $\mathcal{A}$. Moreover, if $E, E' \in \mathcal{A}$ and are disjoint, then repeated use of part (iv) implies $\widehat{\mathcal{A}}\backslash(E \cup E')$ is in $\mathcal{A}$, thus $E \cup E'$ must also be in $\mathcal{A}$. Together, these properties imply that homogeneous systems are λ-systems.

Homogeneous collections are our candidates for the kind of restricted domains on which the decision maker may be probabilistically sophisticated. Consider the following two conditions on $\succcurlyeq$:

Axiom A (*Event Archimedean Property—CS06*). Any set $\mathcal{A} \subseteq \Sigma$ of non-null, pairwise disjoint events, such that $e \approx e'$ for every $e, e' \in \mathcal{A}$, is necessarily finite.

Axiom N (*Event Non-satiation—CS06*). For any pairwise disjoint $E, A, E' \in \Sigma$, if $E \approx E'$ and A is non-null, then no subevent of E' is exchangeable with $E \cup A$.

Axiom N is weaker than the following condition:

Axiom N′ (*Strong event non-satiation*). For any disjoint $E, E', A \in \Sigma$, if $x(E \cup A)x'E'f \sim xEx'(E' \cup A)f$ for every $x, x' \in X$ and $f \in \mathcal{F}$ then A is null.

Lemma 1. *Axiom N′ implies Axiom N.*

CS06 demonstrate that Axioms N and A are sufficient for probabilistic sophistication on any algebraic or finite homogeneous collection. The stronger form of non-satiation, Axiom N′, is therefore only required in the case of an infinite non-algebraic homogeneous collection. [10] We are now ready to state a central result of the paper:

Theorem 1. *Assume Axioms A and N′. Then if $\mathcal{A}$ is a homogeneous collection, there is a unique, solvable, and finitely additive measure, μ, representing $\succcurlyeq^C$ on $\mathcal{A}$. The measure μ is either atomless or purely and uniformly atomic. Moreover, if f and f' induce the same lottery with respect to μ on $\mathcal{A}$, then $f \widehat{\mathcal{A}} g \sim f' \widehat{\mathcal{A}} g$ for every $g \in \mathcal{F}$.*

Theorem 1 confirms that homogeneous collections appropriately depict a basis for 'small worlds' and for addressing the question posed in the Introduction. Because we do not rely on monotonicity or continuity in deriving this result, Theorem 1 also provides a characterization of λ-system probabilistic sophistication that is more parsimonious than others offered in the literature (see [8,40,23]).

Homogeneity is a demanding requirement. In particular, comparability between two events is not sufficient to conclude that they are members of a homogeneous collection. At the same time, every homogeneous collection containing at least two non-null events also contains strict subsets that are homogeneous. For instance, any event along with the empty set forms a homogeneous collection. For a less trivial example, consider that any homogeneous collection with two or more non-null events also contains at least two disjoint events, E and E', such that $E \approx E'$; for any two such events, $\{\emptyset, E, E', E \cup E'\}$ is homogeneous. In characterizing a 'conditional small world' we focus on the following non-trivial construction:

Definition 6. $\mathcal{A} \subseteq \Sigma$ is a **conditional small world event domain** if it is a homogeneous subset of Σ containing more than one non-null event.

For brevity, we will henceforth refer to any conditional small world event domain simply as a 'conditional small world'. The term 'conditional' is appropriate when the universal set of the small world is not the full space. Otherwise, if $\widehat{\mathcal{A}} = \Omega$ then $\mathcal{A}$ can also be referred to as an *unconditional small world*. Finally, we say that $\mathcal{A}$ is *maximal* if for any $\mathcal{A}' \subseteq \Sigma$ containing a non-null event, $\mathcal{A}' \subseteq \mathcal{A}$ whenever $\mathcal{A}' \cup \mathcal{A}$ is homogeneous.

Note the following properties: (i) If Σ is a small world then it is the only maximal conditional small world (up to null sets), and under the hypothesis of Theorem 1 the decision maker exhibits global probabilistic sophistication; (ii) if Σ does not contain a conditional small world, then no two non-null and disjoint events are exchangeable; finally, (iii) if $\mathcal{A}, \mathcal{A}' \subseteq \Sigma$ are conditional small worlds, with $\mathcal{A}' \subseteq \mathcal{A}$ and $\mathcal{A}$ maximal, then the measure $\mu_{\mathcal{A}}$ from Theorem 1 coincides with the corresponding measure $\mu_{\mathcal{A}'}$ when the former is restricted to $\mathcal{A}'$.

3.1. Utility representations

It is useful to investigate conditions that ensure the existence of a continuous utility representation over the act-induced lottery space of a conditional small world. Other approaches (e.g., [33,25,15,8]) impose variants of Savage's P6, [11] and rely on some form of monotonicity

[10] A counter example demonstrating that N and A are not sufficient in the case of an infinite non-algebraic homogeneous collection is non-trivial to construct. We are able to supply such an example upon request for interested readers.

[11] Savage's P6 requires that whenever $f \succ g$, then for any $x \in X$ there is a sufficiently fine finite partition of Ω, say $\{E_i\}_{i=1}^n \subset \Sigma$, such that $xE_i f \succ g$ and $f \succ xE_i g$ for every $i = 1 \ldots n$.

to establish the continuity properties of a utility representation. The absence of monotonicity in our approach prevents us from being able to follow this path because, on its own, P6 is not enough to establish that $\succcurlyeq$ is a dense ordering, a necessary condition for the existence of a utility representation. [12] Instead, we adopt a notion of continuity that is more appropriate for our setting.

Henceforth assume X is a connected and compact metric space, Ω is a separable metric space, and Σ is the Borel σ-algebra of Ω. The *topology of pointwise convergence* on $\mathcal{F}$ is defined to be the relative topology with respect to the product topology on X^{Ω}. Our continuity axiom can now be stated:

CNT (*Continuity*). The sets $\{g \succeq f | g \in \mathcal{F}\}$ and $\{f \succeq g | g \in \mathcal{F}\}$ are closed in the topology of pointwise convergence on $\mathcal{F}$.

Adding this axiom leads to a continuous utility representation over induced lotteries.

Proposition 1. *$\succcurlyeq$ Satisfies CNT if and only if it has a continuous utility representation* (*in the topology of pointwise convergence*).

Let $\overline{\Delta(\mu)}$ denote the closure of $\Delta(\mu)$ in the topology of weak convergence. A real valued function, U, over any $\Delta' \subseteq \Delta$ is said to be non-satiated in probability if, whenever $p, p+q \in (0,1)$, $U\big((p+q)x \oplus (1-p-q)x'\big) = U\big(px \oplus (1-p)x'\big)$ for all $x, x' \in X$ implies $q = 0$.

Theorem 2. *Assume Axioms A, N', CNT, and that $\mathcal{A} \subseteq \Sigma$ is a conditional small world domain. Let μ be the measure that represents $\succcurlyeq^C$ on $\mathcal{A}$ (as guaranteed by Theorem 1). Then for each $h \in \mathcal{F}$ and f, g adapted to $\mathcal{A}$ there exists a function, $U_h : \overline{\Delta(\mu)} \mapsto \mathbb{R}$ that is non-satiated in probability and continuous in the topology of weak convergence, such that*

$$f \stackrel{\frown}{\mathcal{A}} h \succcurlyeq g \stackrel{\frown}{\mathcal{A}} h \Leftrightarrow U_h(L_{\mu,f}) \geqslant U_h(L_{\mu,g}).$$

Theorem 2 provides a representation of the decision maker's risk preference among acts that are adapted to conditional small worlds. Likewise, it is sufficient to characterize the risk preferences of the decision maker within a conditional small world in order to pin down their choice behavior over acts adapted to that small world.

3.2. Examples of conditional small worlds and associated choice representations

3.2.1. The three-color urn

Consider a single draw from a three-color urn containing 30 red balls, plus another 60 balls, each of which can be either blue or green. In addition, to make the state space non-atomic, consider a draw from $x \in [0, 1]$ using some uniform mechanism. A state is identified by the color of the ball drawn and the number drawn from [0, 1]. Refer to the event of drawing a ball of a certain color by the first letter of the color drawn, and let the union of color events be represented by the conjunction of the appropriate letters (e.g., *GB* is the event 'a Green or Blue ball is drawn'). An event is denoted by $C \times I$ where $C \in \mathcal{C} \equiv \{R, G, B, RG, RB, GB, RGB\}$ and I is a Borel set in [0, 1]. As in Ellsberg's [6] original experiment, assume that decision makers are ambiguity

[12] The relation $\succcurlyeq$ is a dense ordering if whenever $f \succ g$, there is some $h \in \mathcal{F}$ such that $f \succ h \succ g$. See Kreps [24] for more details.

averse in the sense that they prefer to bet on $C \times I$ than $C' \times I'$, for $I = I' = [0, 1]$, whenever $C = R$ while C' is either G or B, and whenever $C = GB$ while C' is either RG or RB. It is natural to assume that these preferences will extend to arbitrary non-null Borel sets $I, I' \subset [0, 1]$ with $m(I) = m(I')$ and $m(\cdot)$ the Lebesgue measure. From this, one concludes that events of the form $R \times I$ and $G \times I'$ (or $B \times I'$) are not comparable via $\succeq^C$. We also take it for granted that the following behavior is observed: $C \times I \approx C \times I'$ where $C \in \mathcal{C}$ and $m(I) = m(I')$; and $R \times I \approx GB \times I'$ where $m(I) = 2m(I')$. Moreover, $G \times I$ (or $B \times I$) is not comparable with $GB \times I'$ whenever $m(I) = 2m(I')$. I.e., while a Savagian would invoke the principle of insufficient reason to justify assigning both events equal likelihood, the decision maker views the event $GB \times I'$ as objective, while the same is not true of $G \times I$. Likewise, $G \times I$ is not exchangeable with $B \times I'$ unless $m(I) = m(I') = 1$. To see this, consider that an 'ambiguity averse' agent strictly prefers diversifying across 'ambiguous' sources. For example, if the payoffs, $x, y \in X$, are monetary and $x > y$, then one expects that $0R \times [0, 1]xG \times ([0, 1]\backslash I) \cup B \times Iy \succ 0R \times [0, 1]xG \times [0, 1]y$ (this is similar to the mixture preference assumption of Gilboa and Schmeidler [14]). We list the conditional small worlds implied by $\succcurlyeq$:

(i) $\mathcal{A}_G$ and $\mathcal{A}_B$ are the algebras generated by events of the form $G \times I$ and $B \times I$, respectively. The representing measure is uniform.
(ii) $\mathcal{A}_{PS}$ is the algebra generated by all events of the form $R \times I$ and $G \times I' \cup B \times I''$, with $m(I') = m(I'')$. The representing measure assigns the probabilities $\frac{1}{3}m(I)$ and $\frac{2}{3}m(I')$ to the former and latter events, respectively. Moreover, $\mathcal{A}_{PS}$ is an unconditional small world.
(iii) $\mathcal{A}_{GB}$ is generated by $G \times I$ and $B \times I'$ where $m(I) = m(I') = 1$. The representing measure is purely and uniformly atomic, and contains two atoms.

There is an additional behavioral constraint that can be superimposed on this structure. Specifically, let ℓ_f^G and ℓ_f^B be the lotteries induced by f on, respectively, $\mathcal{A}_G$ and $\mathcal{A}_B$. Likewise, define ℓ_f^R to be the lottery $\{(m(f^{-1}(x) \cap R \times [0, 1]), x)|x \in X\}$. For any $f, g \in \mathcal{F}$, if $\ell_f^G = \ell_g^B$ and $\ell_f^B = \ell_g^G$, while f and g agree on $R \times [0, 1]$, then it appears intuitively sensible that the decision maker will be indifferent between f and g.

To construct a continuous utility representation exhibiting this small world structure and satisfying Theorem 2, consider that every act is adapted to $\mathcal{A}_G$ and $\mathcal{A}_B$. Suppose we make the assumption that one can replace $f = fRfGfB$ with $fRfG\hat{c}_B B$, using some certainty equivalent, $\hat{c}$, that is independent of how f assigns outcomes to G. [13] Then symmetry implies a representation of the form

$$V(f) = \Phi\Big(\frac{1}{3}\ell_f^R \oplus \frac{2}{3}\phi\big(\hat{c}(\ell_f^G), \hat{c}(\ell_f^B)\big)\Big), \tag{1}$$

where Φ represents $\succcurlyeq$ over $\mathcal{A}_{PS}$, and the symmetric function $\phi : \mathbb{R} \times \mathbb{R} \mapsto \mathbb{R}$ is the certainty equivalent for the conditional small world $\mathcal{A}_{GB}$ (spanned by two atoms). Each of ϕ, $\hat{c}(\ell_f^G)$ and $\hat{c}(\ell_f^B)$ can depend on ℓ_f^R, but we suppress this for notational convenience. If f is an act such that $\ell_f^G = \ell_f^B \equiv \ell_f^{GB}$, then it is simultaneously adapted to $\mathcal{A}_{PS}$, $\mathcal{A}_G$, and $\mathcal{A}_B$. This imposes the following relationship between $\Phi(\cdot)$ and $\hat{c}$,

$$\Phi\Big(\frac{1}{3}\ell_f^R \oplus \frac{2}{3}\hat{c}(\ell_f^{GB})\Big) = \Phi\Big(\frac{1}{3}\ell_f^R \oplus \frac{2}{3}\ell_f^{GB}\Big). \tag{2}$$

[13] We make no attempt to justify this 'small world separability' assumption. This and other issues, such as state dependence, are further discussed in the 2003 working paper version of this paper (available from the authors upon request).

Thus the specification of Φ completely pins down $\hat{c}$. For example, if $\Phi(L) = \frac{\int w(x)u(x)\,dL(x)}{\int w(x)\,dL(x)}$ corresponds to Chew's [3] weighted utility model, and $L = \frac{1}{3}\ell_1 \oplus \frac{2}{3}\ell_2$, then $\hat{c}(\ell_2)$ solves

$$\frac{\frac{1}{3}\int w(x)u(x)\,d\ell_1(x) + \frac{2}{3}\int w(x)u(x)\,d\ell_2(x)}{\frac{1}{3}\int w(x)\,d\ell_1(x) + \frac{2}{3}\int w(x)\,d\ell_2(x)} = \frac{\frac{1}{3}\int w(x)u(x)\,d\ell_1(x) + \frac{2}{3}u\big(\hat{c}(\ell_2)\big)w\big(\hat{c}(\ell_2)\big)}{\frac{1}{3}\int w(x)\,d\ell_1(x) + \frac{2}{3}w\big(\hat{c}(\ell_2)\big)}.$$

Assuming monotonicity, Ellsbergian attitudes (i.e., 'ambiguity aversion') are further implied by $\phi(x, y) < \hat{c}\big(\frac{1}{2}x \oplus \frac{1}{2}y\big)$ for $x \neq y$. Thus using conditional small worlds to model behavior that is consistent with both Allais' [1] and Ellsberg's [6] 'paradoxes' is straight forward.

3.2.2. *The two-color urn*

Consider an experiment in which two balls are drawn independently from two different urns: one containing an equal number of red and black balls, while the other contains an indefinite mixture of red and black balls. In order to work with a non-atomic state space, we once again posit an auxiliary draw, $x \in [0, 1]$, using some uniform mechanism. Let R_k refer to the draw of red from the known urn, $R_k B_u$ refer to the simultaneous draw of red from the known urn and black from the unknown urn, and define B_k, R_u, B_u, $R_k R_u$, $B_k B_u$, $B_k R_u$ similarly. The event $\{R_u B_k\} \times I$, where I is a Borel set of [0, 1], has the obvious interpretation and the event space algebra is generated by all events of this form. For notational ease, identify $\{R_u\}$ with $\{R_u B_k, R_u R_k\}$, etc. Experiments and casual introspection suggest that $R_k R_u \times I \approx B_k R_u \times I'$ and $R_k B_u \times I \approx B_k B_u \times I'$ whenever $m(I) = m(I')$. In addition, a preference for diversifying good outcomes over ambiguous events (as in the three-color urn example) implies that neither of $R_k R_u \times I$ nor $B_k R_u \times I$ is exchangeable with any of $R_k B_u \times I'$ or $B_k B_u \times I''$. For a similar reason, $B_u \times I$ is not exchangeable with any $R_u \times I'$ unless $m(I) = m(I') = 1$. The implied conditional small world event domains follow:

(i) $\mathcal{A}_{PS}$ is the algebra generated by all events of the form $R_k \times I$ and $B_k \times I'$. The representing measure is uniform. Moreover, $\mathcal{A}_{PS}$ is an unconditional small world.
(ii) $\mathcal{A}_{\mathrm{R}_u}$ (resp. $\mathcal{A}_{B_u}$) is the algebra generated by all events of the form $R_k R_u \times I$ and $B_k R_u \times I'$ (resp. $R_k B_u \times I$ and $B_k B_u \times I'$). The representing measure is uniform.
(iii) $\mathcal{A}_u$ is the algebra generated by $R_u \times I$ and $B_u \times I'$, where $m(I) = m(I') = 1$. The representing measure is purely and uniformly atomic, and contains two atoms. Note that this too is an unconditional small world.

The two-color urn is analogous to the three-color urn conditional on events in $G \cup B$. In the latter case, the ratio of green to blue balls is uncertain, as is the ratio of red to black balls in the two-color urn. One can therefore identify $\mathcal{A}_G$ with $\mathcal{A}_{\mathrm{R}_u}$ and $\mathcal{A}_B$ with $\mathcal{A}_{B_u}$. Letting $\ell_f^{R_u}$ (resp. $\ell_f^{B_u}$) be the lottery induced by f on $\mathcal{A}_{\mathrm{R}_u}$ (resp. $\mathcal{A}_{B_u}$), then, as in the three-color urn example, an additional behavioral constraint is that for any $f, g \in \mathcal{F}$, if $\ell_f^{R_u} = \ell_g^{B_u}$ and $\ell_g^{R_u} = \ell_f^{B_u}$ then $f \sim g$. In analogy with Eq. (1), a representation satisfying Theorem 2 and first order dominance might take the form,

$$V(f) = \phi\Big(\hat{c}(\ell_f^{R_u}), \hat{c}(\ell_f^{B_u})\Big),$$

where $\hat{c}(L)$ is a certainty equivalent functional satisfying first order dominance, while $\phi : \mathbb{R} \times \mathbb{R} \mapsto \mathbb{R}$ is increasing, symmetric, continuous, $\phi(x, x) = x$, and $\phi(x, y) < \hat{c}\big(\frac{1}{2}x \oplus \frac{1}{2}y\big)$.

Preferences over bets with objective probabilities are represented by $\hat{c}$, while ϕ captures ambiguity attitudes.

3.2.3. Two temperature scales

Consider a state space generated by the realization of future temperature in New York City and Lhasa, normalized so as to lie in $[0,1]^2$. An event is a product, $I \times I'$ where each of I and I' is a Borel set of [0, 1]. Assuming that the decision maker has diversification preferences over temperature intervals in Lhasa but not in New York City (i.e., the decision maker is ambiguity averse with respect to Lhasa's temperature), it is sensible to assume the following conditional small world structure:

(i) $\mathcal{A}_{n,l}$ is the algebra generated by NY temperature events of the form $I_n \times \{l\}$ where I_n is a Borel subset of [0, 1], and $l \in [0,1]$ is the temperature in Lhasa.
(ii) $\mathcal{A}_n$ is the algebra generated by events of the form $I \times [0,1]$. Thus $\mathcal{A}_n$ is an unconditional small world.
(iii) $\mathcal{A}_l$ is the algebra generated by events of the form $[0,1] \times I$. Thus $\mathcal{A}_l$ is also an unconditional small world.

A further intuitive behavioral constraint is that the representing measure on $\mathcal{A}_{n,l}$ is the same as that for $\mathcal{A}_{n,l'}$ for any $l, l' \in [0,1]$ and that, for any act $f \in \mathcal{F}$, the decision maker is indifferent to permuting the distribution of payoffs awarded by f on $[0,1] \times \{l\}$ with that awarded by f on $[0,1] \times \{l'\}$.

In both urn-based examples the domain corresponding to 'ambiguity' corresponds to an equal and binary partition of the state space. In the temperature example, this structural limitation is relaxed because the 'ambiguous domain' (i.e., temperature intervals in Lhasa) is richer. Let $\ell_f^{n,l}$ correspond to the lottery induced by f on the domain $\mathcal{A}_{n,l}$. Then the following representation satisfies Theorem 2,

$$V(f) = c_u\big(\hat{L}_f\big),$$

where $c_u(L)$, is a certainty equivalent functional over lotteries, $\hat{L}_f$ is the lottery corresponding to the distribution function

$$F(x) \equiv m\left(\left\{\hat{c}(\ell_f^{n,l}) \leqslant x | l \in [0,1]\right\}\right),$$

m is the Lebesgue measure on [0, 1], and $\hat{c}(\ell_f^{n,l})$ is a certainty equivalent functional over lotteries (reflecting risk-preferences on $\mathcal{A}_{n,l}$). Ambiguity aversion can be captured by requiring $c_u(L) < \hat{c}(L)$ for any non-degenerate lottery.

A particular instance of such a representation is

$$V(f) = \int_0^1 v \circ u^{-1}\left(\int_0^1 u(t_n, t_l)\, d\mu_n(t_n)\right) d\mu_l(t_l),$$

where the real-valued functions, u and v, defined over [0, 1], are increasing and v is strictly more concave than u. This representation coincides with that developed in Ergin and Gul [9] and Nau [27], and corresponds to a type of two-stage or recursive utility representation for modeling ambiguity aversion outside of the standard multiple-prior setting. A necessary and sufficient condition for this representation to hold in our setting is that risk preferences be separable (i.e., satisfy Savage's Sure Thing Principle) within each small world.

4. Application: subjectively unambiguous events

One strand of the literature on decision making under uncertainty attempts to distinguish events on which a decision maker is probabilistically sophisticated, from those on which the decision maker's likelihood relation does not admit a probabilistic representation. The former are often called 'unambiguous' while the latter are termed 'ambiguous' events. For instance, Epstein [7] asserts that there are exogenously given unambiguous events on which the decision maker is probabilistically sophisticated in the Machina-Schmeidler sense and defines the decision maker's risk attitude there as being ambiguity neutral, thus providing a benchmark for ambiguity aversion.

In their 2001 paper, Epstein and Zhang posited four desiderata for a definition of subjectively unambiguous events: (D1) Behavioral; (D2) model free; (D3) explicit and constructive; and (D4) consistent with probabilistic sophistication on unambiguous acts. An event, $B \in \Sigma$, is subjectively unambiguous in their sense (henceforth **EZ-unambiguous**) if the following two conditions are satisfied (i) for any $x, x', z, z' \in X$, $f \in \mathcal{F}$ and $E, E' \in \Sigma$ such that $E, E' \subseteq B^c$ and $E \cap E' = \emptyset$,

$$zBxEx'E'f \succeq zBx'ExE'f \Rightarrow z'BxEx'E'f \succeq z'Bx'ExE'f.$$

and, (ii) condition (i) holds if B is everywhere replaced by B^c. Epstein and Zhang [8] show that the set of subjectively unambiguous events is a λ-system under Savage's P5 and some technical conditions, and further, that desideratum D4 is satisfied under P3 and modified forms of P4 and P6.

In our setting, an attempt can also be made to identify subjectively unambiguous events in relation to homogeneous collections of events. Specifically, define $E \in \Sigma$ to be **EB-unambiguous** (for 'exchangeability-based unambiguous') if it belongs to *some* maximal unconditional small world, $\mathcal{A} \subseteq \Sigma$ (i.e., $\widehat{\mathcal{A}} = \Omega$). This definition satisfies the Epstein–Zhang desiderata in that it is: (i) behavioral by virtue of being exchangeability-based, (ii) constructive in the sense that one can in principle systematically investigate whether an unconditional small world can constructed from the event under scrutiny and its complement, (iii) consistent with probabilistic sophistication by virtue of Theorem 1, and (iv) does not refer to the decision maker's ambiguity or risk attitudes, or to exogenous structure about the state space.

Given that our definition appears to satisfy the Epstein–Zhang desiderata, it is appropriate to make a more careful comparison between the two. We first note that the two definitions are not nested in that events classified as unambiguous by Epstein and Zhang [8] may not be so classified by us, and vice versa. The latter is easy to see from the two urn example where our definition regards any single-urn event as unambiguous since the decision maker is probabilistically sophisticated within each urn small world. To see the former, consider the following example:

Example 1. A randomly thrown dart can land in one of four rectangular regions: A, B_1, B_2, and C. The following information is available: the regions represented by A and C are each at least as large as either of the regions B_1 or B_2; moreover, A is larger than C if and only if B_1 is larger than B_2. Let Ω be $[0, 1] \times \{A, B_1, B_2, C\}$, so that events partition each of A, B_1, B_2, and C. [14] Furthermore, let $X \equiv [-1, 1]$, let $E_A[f]$ and $\text{var}_A[f]$ be, respectively, the expected value and variance of the act f conditional on the dart falling in region A (using a uniform distribution),

[14] E.g., one can speak of the dart landing in the top third of region A. We identify the event $[0, 1] \times A$ with A, and similar for the other regions.

and similar for B_1, B_2, and C. Finally, define $U_A[f] \equiv E_A[f] - \frac{1}{8}\text{var}_A[f]$, and similar for other regions.

Consider the utility representation that assigns the act f a utility of

$$V(f) = \frac{1}{3}\left(U_A[f] + U_C[f] + \frac{U_{B_1}[f] + U_{B_2}[f]}{2}\right) \times \min\left\{1, 1 - \frac{(E_A[f] - E_C[f])(E_{B_1}[f] - E_{B_2}[f])}{16}\right\}.$$

This can be interpreted as follows: The decision maker first calculates a mean-variance utility for each of the regions, and these are then averaged by assigning each of A and C a probability of $\frac{1}{3}$, and a probability of $\frac{1}{6}$ to each of B_1 and B_2; she then assesses utility by adjusting this calculation by a 'hedging premium' that gives preference to acts in which a higher (resp. lower) average payoff on A versus C is balanced by a higher (resp. lower) average payoff on B_2 versus B_1. It is straight forward to check that the representation is monotonic. It should also be clear that $A \approx B_1 \cup B_2$ and $C \approx B_1 \cup B_2$. On the other hand, it is certainly not the case that $A \approx C$ due to the 'preference for hedging' (i.e., the asymmetry between payoffs on A and C arising in the last term of the utility function).

It is straight forward to show that the set of EZ-unambiguous events contains $B_1 \cup B_2$ and $A \cup C$. On the other hand, $B_1 \cup B_2$ cannot be part of a homogeneous collection with envelope Ω, because the symmetry of the utility function would then counterfactually require that $A \cup C$ contains two exchangeable subevents such that each is also exchangeable with $B_1 \cup B_2$.

Our definition of unambiguity is 'small world' compatible in the sense that one does not require knowledge about the decision maker's attitudes toward events *outside* of the homogeneous system used to establish unambiguity. This is not the case in the Epstein–Zhang definition, which involves assessing *all* pairs of events E' and E'' in the complement of E. This can pose difficulties if 'big world' events (especially those not modeled) might spoil the unambiguity of what might otherwise appear to be clearly unambiguous. To see how such a concern might arise, return once more to Example 1.

Note first that the unique unconditional small world with envelope Ω—and therefore the unique λ-system of EB-unambiguous events—consists of all events containing the *same proportion* of each region (e.g., 'the dart falls in the top half of one of the regions'); we refer to this small world as $\mathcal{A}_{[0,1]}$, since it is essentially equivalent to the algebra of subsets of the unit interval. If $E \in \mathcal{A}_{[0,1]}$ then the event generated by permuting the identity of the four regions is also in $\mathcal{A}_{[0,1]}$. The decision maker is probabilistically sophisticated and has monotonic mean-variance preferences with respect to all acts adapted to $\mathcal{A}_{[0,1]}$. On the other hand, no non-null event in $\mathcal{A}_{[0,1]}$ is EZ-unambiguous. To see this, consider without loss of generality the event $E = I \times A \cup B_1 \cup B_2 \cup C$ where $I \subset [0, 1]$ has Lebesgue measure of $0 < q < 1$. Let I^c be the complement of I in [0, 1], $E' = I^c \times A$, $E'' = I^c \times C$, $\hat{B}_1 \equiv I^c \times B_1$, and $\hat{B}_2 \equiv I^c \times B_2$. Consider acts of the form $zExE'x'E''y_1\hat{B}_1y_2\hat{B}_2$, where $z, x, x', y_1, y_2 \in [-1, 1]$, while $y_1 \neq y_2$ and $x \neq x'$. Then a necessary condition for E to be EZ-unambiguous is that the expression

$$qz + \frac{1-q}{3}\left(x + x' + \frac{y_1 + y_2}{2} - \frac{q}{8}\left((x - z)^2 + (x' - z)^2 + \frac{1}{2}(y_1 - z)^2 + \frac{1}{2}(y_2 - z)^2\right)\right)$$

does not change sign as z varies in $[-1, 1]$. To see that this cannot be, set x, x', y_1 and y_2 to be arbitrarily close to zero and vary z from -1 to 1. To summarize, whereas the unambiguity of E in our setting depends only on events in the unconditional small world, $\mathcal{A}_{[0,1]}$, the Epstein–Zhang definition explicitly depends on examining events outside of $\mathcal{A}_{[0,1]}$ (e.g., the events E' and E'' in the example). In particular, one might worry that given any seemingly unambiguous collections of events, such as an algebra on the unit interval with the uniform measure, it might be possible to make events in the collection EZ-*ambiguous* by taking account of choice attitudes in an enriched space.

When Σ is an atomless σ-algebra, we may also alternatively distinguish unambiguous events as follows. First, let $\{a_{i,m}\}_{i=1}^{m}$ be a *uniform m-partition of* Ω whenever the $a_{i,m}$'s are pairwise disjoint, $\bigcup_{i=1}^{m} a_{i,m} = \Omega$, and $a_{i,m} \approx a_{j,m}$ for any $i, j = 1, \ldots, m$. Without loss of generality set $a_{0,m} \equiv \emptyset$. Define $E \in \Sigma$ to be subjectively EB-unambiguous if it is equal to Ω, is null, or whenever for every $m \geqslant 1$ there exists a uniform m-partitions of Ω and $0 \leqslant j \leqslant m$ such that

$$\bigcup_{i=0}^{j} a_{i,m} \subset E \subseteq \bigcup_{i=0}^{j+1} a_{i,m}.$$

Here, an event is said to be EB-unambiguous if it can be approximated arbitrarily well by a union of events from a partition of mutually exchangeable events. This alternative definition also captures the intuition underlying the frequentist view of objective probability and likewise satisfies the Epstein–Zhang desiderata.

Our setting admits the possibility that there are more than one maximal and unconditional small world, in which case no single small world spans the set of unambiguous events. Here, an event, which appears intuitively to be an ambiguous event, may be classified as unambiguous by our definition. For instance, in the case of a decision maker facing Ellsberg's two urns, an outside observer without prior knowledge about which urn contains an equal number of balls would not be able to discern the 'known' versus 'unknown' urn based on our definition alone. What is needed is additional information concerning the decision maker's preference for betting on events with known probabilities versus those without.

Recent literature does not appear to offer a consensus on the definition of unambiguous events (see, for instance, the views expressed in Wakker [39]; and Nehring [30]). As with our case, the Epstein–Zhang definition may give a confounded analysis of ambiguous events in the two urn example. The definition of subjective ambiguity in Ghirardato and Marinacci [12] and Ghirardato, Marinacci and Maccheroni [13], on the other hand, rely on prior knowledge of ambiguity attitudes, while Epstein [7], Nau [27], and Klibanoff, Marinacci and Mukerji [21] appeal to the existence of an exogenously specified set of unambiguous events on which the decision maker is probabilistically sophisticated.

One advantage of the definition of unambiguous events we offer is that it is clear when it should be supplemented with additional information in order to conform to intuition: I.e., whenever there are multiple unconditional and maximal small worlds. More generally, from the perspective of the decision maker, the presence of distinct risk attitudes over endogenously defined small worlds can be interpreted as signaling the existence of 'ambiguity' or 'Knightian uncertainty'. Correspondingly, being ambiguity averse toward uncertainty arising from one source with respect to another may be interpreted in terms of comparative risk aversion. This is, after all, what the Ellsberg Paradox reveals: two distinct certainty equivalents for two 50–50 bets.

5. Conclusion

The literature on probabilistic sophistication concerns contingencies that Savage [33] envisaged as states of the 'big' world—a description of the decision maker's world leaving no relevant aspect untouched. At the same time, Savage observed that decisions are generally made in smaller worlds, which contain events summarizing the relevant aspects of the contingencies pertaining to specific decision situations. This perspective leaves us the question of consistency in decision making from one small world to another. Under Savage's formulation, or more generally probabilistic sophistication, this question has a ready answer: Events in any small world are comparable (in our language) to events in any other small world. In the aggregate, small worlds are all derived from and remain similar to a single big world.

In this paper, we attempt a precise definition of decision making involving small worlds that may be distinct. Drawing on the de Finetti–Ramsey idea of exchangeability, we develop a concept of comparability to capture the intuition of similarity among events, which differ from each other only by a sense of likelihood. This leads to our definition of a small world event domain as a collection of comparable events. When conditioning on events within such a domain, the latter can be viewed as an endogenously induced small world. This paper offers an efficient and weak axiomatization of probabilistic sophistication on (possibly) multiple small worlds. Moreover, we illustrate our approach in several relatively simple settings, including the now classical Ellsberg paradoxes.

There is much room for follow up research in terms of further theoretical development as well as experimental validation of specific small world models. At the same time, we offer simple and empirically testable specifications to model economic behavior involving multiple sources of uncertainty, ranging from behavioral games to markets. Overall, our approach enables decision theoretic modeling of economic settings at the small world level, deferring the question of consistent extension across distinct small worlds as and when the need arises. As Savage puts it, cross the bridge when you get to it.

Acknowledgments

We acknowledge helpful feedback from participants of seminar workshops at UC Berkeley, UBC, Bielefeld, Caltech, Heidelberg, HKUST, INSEAD, NUS, UC Irvine, and UCLA, and participants of RUD 2003 in Milan and FUR 2004 in Paris. We particularly benefited from the comments of Eddie Dekel, Paolo Ghirardato, Itzhak Gilboa, Simon Grant, Mark Machina, Marzena Rostek, Uzi Segal, Costis Skiadas, Sun Yeneng, and especially, Peter Wakker. Support from the Research Grants Council of Hong Kong (HKUST-6304/03H) is gratefully acknowledged.

Appendix A.

Proof of Lemma 1. Suppose $E, A, E' \in \Sigma$ are disjoint, such that $E \approx E'$, and A is non-null. Assume it is not the case that $E \cup A \succ^C E'$. Since $E \cup A$ and E' are comparable, $E' \succeq^C E \cup A$. Thus E' contains a subset, ξ', that is exchangeable with $E \cup A$. In particular, by exchanging ξ' for $E \cup A$, we have for any $x, x' \in X$ and $f \in \mathcal{F}$ that

$$x'(E \cup A)xE'f \sim x'\xi'x\big((E \cup A) \cup (E' \backslash \xi')\big)f = x'\xi'x\big(E \cup (A \cup E' \backslash \xi')\big)f.$$

Similarly, by exchanging E with E' it follows that:

$$x'(E \cup A)xE'f \sim x'(E' \cup A)xEf = x'\big(\xi' \cup (A \cup E' \backslash \xi')\big)xEf.$$

Note that the set $A \cup E' \backslash \xi'$ is not null, since A is not null. Axiom N′ is therefore violated, meaning that it cannot be the case that $E' \succeq^C E \cup A$. Thus $E \cup A \succ^C E'$. □

Proof of Theorem 1. Begin with several useful results and constructions:

Lemma A.1. *Axiom N′ implies for any disjoint $E, E' \in \Sigma$: $E \sim^C E' \Leftrightarrow E \approx E'$.*

Proof. $E \sim^C E' \Rightarrow \exists e \subseteq E$ with $e \approx E'$. $E \backslash e$ must be null (in which case $E \approx E'$). Otherwise, $E' \sim^C e \cup (E \backslash e)$ implies that E' contains a subevent exchangeable with $e \cup (E \backslash e)$ in violation of N (and therefore N′). Now, $E \approx E'$ implies $E \succeq^C E'$ and $E' \succeq^C E$, thus implying $E \sim^C E'$. □

Lemma A.2. *If $a, b, c \in \mathcal{A}$, $b \cap c = \emptyset$, $b \succ^C c$, and $a \sim^C b$ then $a \succ^C c$.*

Proof. Let $a_b \equiv a \cap b$, $a_c \equiv a \cap c$, $a_0 \equiv a \backslash (b \cup c)$, $b_0 \equiv b \backslash a$ and $c_0 \equiv c \backslash a$. Assume that $c \succcurlyeq^C a$. Then there is some event $c' \subseteq c_0$ with $a_0 \cup a_b \approx c'$. From $b \succ^C c$, there is some event $b' \subset b$ with $b' \approx c$ and $b \backslash b'$ not null. For any $x, y \in X$ and $f \in \mathcal{F}$, consider the act $g \equiv x\Big(b \cup (c_0 \backslash c')\Big)y(a_0 \cup a_c \cup c')f$. Using the exchangeability of $a_0 \cup a_c$ and b_0 (Lemma A.1), $g \sim x\big((a \cup c) \backslash c'\big)y(b_0 \cup c')f$. The exchangeability of c' and $a_0 \cup a_b$ can be used to now give, $g \sim xcy(b \cup a_0 \cup a_b)f$. Finally, using the exchangeability of b' and c gives, $g \sim xb'y(c \cup (a \cup b) \backslash b')f$. Setting $A \equiv (b \backslash b') \cup (c_0 \backslash c')$ and applying Axiom N′ implies A is null, a contradiction. Thus $a \succ^C c$. □

Lemma A.3. *If $a, b, c \in \mathcal{A}$, $b \cap c = \emptyset$, $b \approx c$, and $a \sim^C b$, then $a \sim^C c$.*

Proof. Define a_0, a_b, a_c, b_0 and c_0 as in the proof of the previous Lemma. For any $x, y \in X$ and $f \in \mathcal{F}$, set $g \equiv xay(b_0 \cup c_0)f$. Using $a_0 \cup a_c \approx b_0$ gives $g \sim xby(a_0 \cup c)f$, and using $b \approx c$ further gives $g = xay(b_0 \cup c_0)f \sim xcy(a_0 \cup b)f$. Either of $a \succ^C c$ or $c \succ^C a$ leads to a contradiction with Axiom N′, thus comparability of a and c through $\succcurlyeq^C$ necessitates $a \sim^C c$. □

Lemma A.4. *If $a_1, a_2, b_1, b_2 \in \mathcal{A}$, $a_1 \cap a_2 = b_1 \cap b_2 = \emptyset$, $a_1 \sim^C b_1$, and $a_2 \sim^C b_2$, then $a_1 \cup a_2 \sim^C b_1 \cup b_2$.*

Proof. Assume otherwise and set $a \equiv a_1 \cup a_2$, $b \equiv b_1 \cup b_2$, $a_0 \equiv a \backslash b$, and $b_0 \equiv b \backslash a$. Then without loss of generality one can assume that $a \succ^C b$ and there are events $a'_1 \subseteq a_1 \backslash b$ and $a'_2 \subseteq a_2 \backslash b$ such that, $a' \equiv a'_1 \cup a'_2$ is not null and $a \backslash a' \sim^C b$. For any $x, y \in X$ and $f \in \mathcal{F}$, let $g = x\big(a' \cup b_0 \cup (b_1 \cap a)\big)y\big((a_0 \backslash a') \cup (b_2 \cap a)\big)f$. Using $a_0 \backslash a' \approx b_0$, one derives $g \sim x\big(a_0 \cup (b_1 \cap a)\big)y\big(b_0 \cup (b_2 \cap a)\big)f$. Now use $a_2 \sim^C b_2$ to deduce that $g \sim x\big(a_1 \cup (b_2 \backslash a)\big)y\big(a_2 \cup (b_1 \backslash a)\big)f$. Finally, use $a_1 \sim^C b_1$ to get $g \sim x\big(b_0 \cup (b_1 \cap a)\big)y\big(a_0 \cup (b_2 \cap a)\big)f$. By Axiom N′, a' is null—a contradiction. Thus $a \sim^C b$. □

Lemma A.5. *Let $A, A', E, E' \in \mathcal{A}$ with $E \subseteq A$, $E' \subseteq A'$, $A \sim^C A'$, $E \sim^C E'$. Then $A \backslash E \sim^C A' \backslash E'$.*

Proof. Let $\hat{E} \equiv A\backslash E$, $\hat{E}' \equiv A'\backslash E'$. For any $x, y \in X$ and $f \in \mathcal{F}$, set $g \equiv x\big((\hat{E}\backslash\hat{E}') \cup (E\backslash A')\big)y\big((\hat{E}'\backslash\hat{E}) \cup (E'\backslash A)\big)f = x\big((A\backslash A') \cup (\hat{E} \cap E')\big)y\big((A'\backslash A) \cup (\hat{E}' \cap E)\big)f$. Using $A \sim^C A'$ gives $g \sim x\big((A'\backslash A) \cup (\hat{E} \cap E')\big)y\big((A\backslash A') \cup (\hat{E}' \cap E)\big)f$. Now using $E \sim^C E'$ gives $g \sim x\big((\hat{E}'\backslash\hat{E}) \cup (E\backslash A')\big)y\big((\hat{E}\backslash\hat{E}') \cup (E'\backslash A)\big)f$. Axiom N′ is not consistent with the previous relation if $\hat{E} \succ^C \hat{E}'$ or $\hat{E}' \succ^C \hat{E}$. Thus, because $\hat{E}$ and $\hat{E}'$ are comparable, it must be that $\hat{E} \sim^C \hat{E}'$. □

The following construction is instrumental:

Definition A.1. For any non-null $E, A \in \Sigma$ such that $E \subseteq A$, a *division of A by E* is a partition of A into events in Σ such that,

(i) E is contained in the partition,
(ii) all but, possibly, one of the partition elements are exchangeable with E,
(iii) the *remainder* event, E_{R}, is non-null and $E \succ^C E_{\mathrm{R}}$.

While a division of A by E may not exist for arbitrary events where $E \subseteq A$, if A and E are part of a small world event domain and Axiom A holds, then such a division can be constructed by definition.

Lemma A.6. *Let $A, A', E, E' \in \mathcal{A}$ with $E \subseteq A$, $E' \subseteq A'$, $A \sim^C A'$, $E \sim^C E'$. Suppose, further, that $\mathcal{D} \subseteq \mathcal{A}$ is a division of A by E and $\mathcal{D}' \subseteq \mathcal{A}$ is a division of A' by E'. Then $\mathcal{D}$ has the same cardinality as $\mathcal{D}'$, if e and e' are non-remainder events then $e \sim^C e'$, and if E_{R} and E'_{R} are remainder events then $E_{\mathrm{R}} \sim^C E'_{\mathrm{R}}$.*

Proof. First, Lemma A.3 implies that if a and b are non-remainder events in $\mathcal{D} \cup \mathcal{D}'$, then $a \sim^C b$. Suppose that $\mathcal{D}$ has n non-remainder elements, denoted by E_i, $i = 1, \ldots, n$, while $\mathcal{D}'$ has n' non-remainder elements (similarly denoted). Lemma A.4 can subsequently be used to show that $\bigcup_{i=1}^{m} E_i \sim^C \bigcup_{i=1}^{m} E'_i$ for any m smaller or equal to $\min(n, n')$. If $n \neq n'$, then suppose without loss of generality that $n < n'$. Lemma A.5 implies that $E_{\mathrm{R}} \sim^C E'_{\mathrm{R}} \cup \bigcup_{j=n+1}^{n'} E'_j$. Because $E \succ^C E_{\mathrm{R}}$ and $E'_{n+1} \sim^C E$, Lemma A.2 implies that $E'_{n+1} \succ^C E_{\mathrm{R}}$—a contradiction. Thus $n = n'$ and Lemma A.5 implies that $E'_{\mathrm{R}} \sim^C E_{\mathrm{R}}$. □

Definition A.2. Suppose $E, A \in \mathcal{A}$ with $E \subseteq A$ and let $\mathcal{D} \subseteq \mathcal{A}$ be a division of A by E. If $\mathcal{D}$ contains a remainder event, E_{R}, then a *first-order refinement of* $\mathcal{D}$ is a union, $\bigcup_{E_\alpha \in \mathcal{D}} \mathcal{D}_\alpha$, where $\mathcal{D}_\alpha \subseteq \mathcal{A}$ is a division of $E_\alpha \cup E_{\mathrm{R}}$ by E_{R}. If $\mathcal{D}$ does not contain a remainder event, then the first-order refinement of $\mathcal{D}$ is $\mathcal{D}$ itself.

Fig. 1 illustrates a first-order refinement. The drawn region represents the event A. The events divided by bold vertical lines represent a division, $\mathcal{D}$, of A by the event E. The hatched event, E_{R}, is the remainder. To obtain the first-order refinement, each of the events, $E, E_2, \ldots, E_5$, is further divided by E_{R}. An argument similar to that used in Lemma A.6 leads to the conclusion that the first-order refinement is a partition of A that consists of at most two types of events: those exchangeable with E_{R} and those exchangeable with the event, $E_{\mathrm{R},1}$, which is the remainder of $E \cup E_{\mathrm{R}}$ by E_{R}. The second type of event is illustrated by any of the five narrow unhatched vertical regions in the figure. Events of the type $E_{\mathrm{R},1}$ can be viewed as remainder

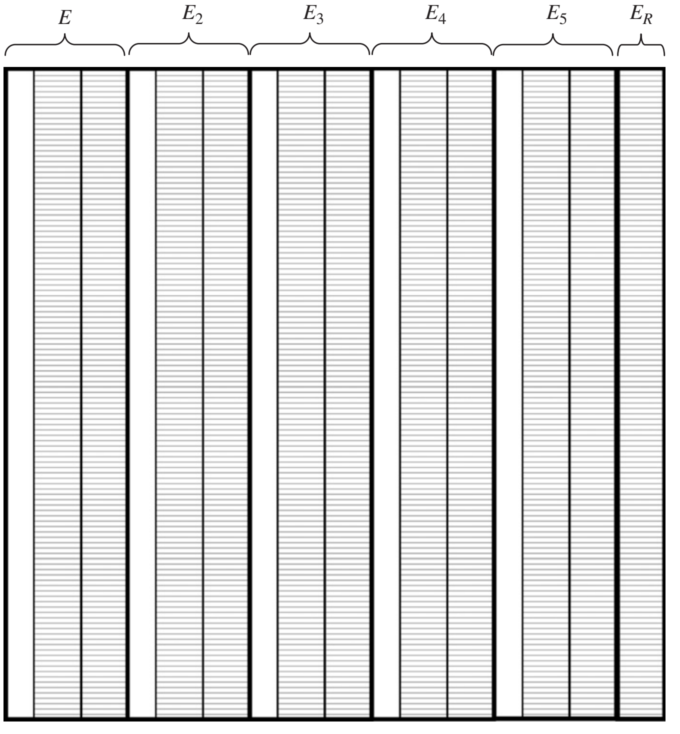

Fig. 1. An illustration of a first-order refinement.

events of the first-order refinement of $\mathcal{D}$. Note, however, that there is more than one such remainder event in a refinement. This construction allows one to define higher order refinements recursively.

Definition A.3. Suppose $E, A \in \mathcal{A}$ with $E \subseteq A$ and let $\mathcal{D} \subseteq \mathcal{A}$ be a division of A by E. Define $\mathcal{D}_0 \equiv \mathcal{D}$ to be a 0th-order refinement of $\mathcal{D}$. For any $n > 0$, $\mathcal{D}_n \subseteq \mathcal{A}$ is an nth-*order refinement of* $\mathcal{D}$ if it is a union, $\bigcup_{E_\alpha \in \mathcal{D}_{n-1}} \mathcal{D}_\alpha$, where $\mathcal{D}_{n-1}$ is an $n-1$th-order refinement of $\mathcal{D}$, $\mathcal{D}_\alpha$ is a division of $E_\alpha \cup E_{\mathrm{R},n-1}$ by $E_{\mathrm{R},n-1}$, and $E_{\mathrm{R},n-1}$ is one of the remainder events in $\mathcal{D}_{n-1}$. If there are no remainder events in $\mathcal{D}_{n-1}$, then $\mathcal{D}_n$ is defined to be $\mathcal{D}_{n-1}$. The nth-order refinement $\mathcal{D}_n$ is said to be a first-order refinement of $\mathcal{D}_{n-1}$.

A sequence, $\{\mathcal{D}_n\}_{n=0}^{\infty}$, of nth-order refinements of some division, $\mathcal{D} \subseteq \mathcal{A}$, and in which $\mathcal{D}_{n+1}$ is a first-order refinement of $\mathcal{D}_n$, is a sequence of progressively (weakly) finer partitions. In particular, any event in $\mathcal{D}_n$ is a union of events in $\mathcal{D}_{n+1}$.

Lemma A.6 can be extended to nth-order refinements.

Lemma A.7. *Let* $A, A', E, E' \in \mathcal{A}$ *with* $E \subseteq A$, $E' \subseteq A'$, $A \sim^C A'$, $E \sim^C E'$. *Suppose, further, that* $\mathcal{D} \subseteq \mathcal{A}$ *is a division of* A *by* E, $\mathcal{D}' \subseteq \mathcal{A}$ *is a division of* A' *by* E', $\mathcal{D}_n$ *is an* nth-*order refinement of* $\mathcal{D}$, *and* $\mathcal{D}'_n$ *is an* nth-*order refinement of* $\mathcal{D}'$. *Then* $\mathcal{D}_n$ *has the same cardinality as*

$\mathcal{D}'_n$, if $e \in \mathcal{D}_n$ and $e' \in \mathcal{D}'_n$ are non-remainder events then $e \sim^C e'$, and if $E_{\text{R}} \in \mathcal{D}_n$ and $E'_{\text{R}} \in \mathcal{D}'_n$ are remainder events then $E_{\text{R}} \sim^C E'_{\text{R}}$.

Proof. Proceed inductively by assuming the result is true for $\mathcal{D}_n$ and $\mathcal{D}'_n$. The assumption is true for $n = 0$ (Lemma A.6). Let $\mathcal{D}_n = \bigcup_{E_\alpha \in \mathcal{D}_{n-1}} \mathcal{D}_\alpha$ and $\mathcal{D}'_n = \bigcup_{E_\alpha \in \mathcal{D}'_{n-1}} \mathcal{D}'_\alpha$. Then the induction hypothesis implies that for each $\mathcal{D}_\alpha$ in $\mathcal{D}_n$ there is a corresponding $\mathcal{D}'_\alpha$ in $\mathcal{D}'_n$. This guarantees that $\mathcal{D}_n$ and $\mathcal{D}'_n$ have the same cardinality. Lemmas A.3 and A.6 imply the remaining claims. □

Proposition A.1. *Let $E, A \in \mathcal{A}$, $\mathcal{D} \subseteq \mathcal{A}$ be a division of A by E, and $\{\mathcal{D}_n\}_{n=0}^{\infty}$ be a sequence of partitions of A such that $\mathcal{D}_n$ is an nth-order refinement of $\mathcal{D}$. Let $\mathcal{A}_n$ denote the algebra generated by $\mathcal{D}_n$. Then $\bigcup_{n=1}^{\infty} \mathcal{A}_n$ is a homogeneous algebra and there is a unique, solvable, and finitely additive measure, μ_E, that represents $\succcurlyeq^C$ over $\bigcup_{n=1}^{\infty} \mathcal{A}_n$. Moreover, μ is either atomless or purely and uniformly atomic.*

Proof. Consider $e, e' \in \bigcup_{n=1}^{\infty} \mathcal{A}_n$. Then $e \in \mathcal{A}_n$ and $e' \in \mathcal{A}_m$ for some finite m, n. Without loss of generality, assume $m \leqslant n$, so $\mathcal{A}_n$ is finer than $\mathcal{A}_m$ and $e, e' \in \mathcal{A}_n$. Thus, $e \cup e', e \cap e' \in \mathcal{A}_n$ and $\bigcup_{n=1}^{\infty} \mathcal{A}_n$ is an algebra.

Because $e, e' \in \mathcal{A}$ and $\bigcup_{n=1}^{\infty} \mathcal{A}_n$ is an algebra, assume without loss of generality that $e \cap e' = \emptyset$ and that $e \succcurlyeq^C e'$. To establish that $\bigcup_{n=1}^{\infty} \mathcal{A}_n$ is homogeneous it is sufficient to demonstrate that there is some $\hat{e} \in \bigcup_{n=1}^{\infty} \mathcal{A}_n$ such that $\hat{e} \subseteq e$ and $\hat{e} \approx e'$. To this end, begin by expressing e as a union of l_{N} non-remainder events and l_{R} remainder events of $\mathcal{D}_n$, while e' is similarly constructed of k_{N} non-remainder and k_{R} remainder events. In light of Lemma A.6, all non-remainder events are pair-wise exchangeable and the same is true for remainder events. Thus in finding a subevent of e that is exchangeable with e' one can begin by pairing remainder events in e with those in e' and do the same with non-remainder events in e and e'. In other words, the event $\hat{e}$ can be constructed by the union of $\min\{l_{\text{R}}, k_{\text{R}}\}$ remainder events, $\min\{l_{\text{N}}, k_{\text{N}}\}$ non-remainder events, and an event, $\hat{e}_1$, that is exchangeable with some event e'_1: The event e'_1 is e' less $\min\{l_{\text{R}}, k_{\text{R}}\}$ of its remainder events and $\min\{l_{\text{N}}, k_{\text{N}}\}$ of its non-remainder events. In particular, if $l_{\text{N}} \geqslant k_{\text{N}}$ and $l_{\text{R}} \geqslant k_{\text{R}}$ then we are done (because $e \succcurlyeq^C e'$, it cannot be that $l_{\text{N}} < k_{\text{N}}$ and $l_{\text{R}} < k_{\text{R}}$). There are two other cases: (i) $\hat{e}_1$ is a subevent of l_1 non-remainder events in $\mathcal{D}_n$ while e'_1 consists of k_1 remainder events in $\mathcal{D}_n$; and (ii) $\hat{e}_1$ is a subevent of l_1 remainder events in $\mathcal{D}_n$ while e'_1 consists of k_1 non-remainder events in $\mathcal{D}_n$.

Case (i): Denote any one of the (non)-remainder events in $\mathcal{D}_n$ as $E_{\text{R},n}$ ($E_{\text{NR},n}$) and consider the partition $\mathcal{D}_{n+1}$. Let q_1 be the number of non-remainder events resulting from the division of $E_{\text{NR},n}$ by $E_{\text{R},n}$. If $q_1 l_1 \geqslant k_1$ then $\hat{e}_1$ can be constructed from k_1 non-remainder events of $\mathcal{D}_{n+1}$ and we are done. Otherwise, construct $\hat{e}_1$ from $q_1 l_1$ non-remainder events of $\mathcal{D}_{n+1}$ and an event, $\hat{e}_2$, that is a subevent of l_1 remainder events in $\mathcal{D}_{n+1}$, and is exchangeable with some e'_2: The event e'_2 consists of $k_1 - q_1 l_1$ non-remainder events in $\mathcal{D}_{n+1}$. Set $l_2 = l_1$, $k_2 = k_1 - q_1 l_1$, and let q_2 be the number of non-remainder events of resulting from the division of $E_{\text{NR},n+1}$ by $E_{\text{R},n+1}$. To be consistent with $e \succcurlyeq^C e'$, it must be the case that $l_2 > q_2 k_2$. Now construct $\hat{e}_2$ from $q_2 k_2$ non-remainder events of $\mathcal{D}_{n+2}$ and an event, $\hat{e}_3$, that is a subevent of $l_2 - q_2 k_2$ non-remainder events in $\mathcal{D}_{n+2}$, and is exchangeable with some e'_3: The event e'_3 consists of k_2 remainder events in $\mathcal{D}_{n+2}$. Setting $l_3 = l_2 - q_2 k_2 < l_2$ and $k_3 = k_2 < k_1$, the problem is identical to the one faced initially (with $\hat{e}_1$ and e'_1), only the partition elements belong to $\mathcal{D}_{n+2}$ instead of $\mathcal{D}_n$, and l_3, k_3 are strictly smaller than their counterparts. An iteration of this construction must eventually stop since l_1 and k_1 are finite. The iteration can only stop, however, if one achieves the construction of $\hat{e} \subseteq e$ exchangeable with e'.

Case (ii): This corresponds to the second step (with $\hat{e}_2$ and e'_2) in Case (i). One can proceed with the iteration from there.

The construction above establishes that for any $e, e' \in \bigcup_{n=1}^{\infty} \mathcal{A}_n$, there is a subevent of $e \backslash e'$ in $\bigcup_{n=1}^{\infty} \mathcal{A}_n$ that is exchangeable with $e' \backslash e$. Thus $\bigcup_{n=1}^{\infty} \mathcal{A}_n$ is a homogeneous algebra. Given Axioms A and N′, the existence of a unique, solvable and finitely additive representing measure, μ_E, that is either atomless or purely and uniformly atomic, proceeds as in CS06. □

Consider non-null $E, E' \in \mathcal{A}$ and let $\mathcal{D}$ (resp. $\mathcal{D}'$) be a division of $\widehat{\mathcal{A}}$ by E (resp. E'), and $\{\mathcal{D}_n\}_{n=0}^{\infty}$ (resp. $\{\mathcal{D}'_n\}_{n=0}^{\infty}$) be a sequence of partitions of $\widehat{\mathcal{A}}$ where $\mathcal{D}_{n+1}$ (resp. $\mathcal{D}'_{n+1}$) is an nth-order refinement of $\mathcal{D}_n$ (resp. $\mathcal{D}'_n$). Let $\mathcal{A}_n$ (resp. $\mathcal{A}'_n$) denote the algebra generated by $\mathcal{D}_n$ (resp. $\mathcal{D}'_n$), and μ_E (resp. μ'_E) correspond to the unique representing measure of $\succcurlyeq^C$ over $\mathcal{A}_E = \bigcup_{n=1}^{\infty} \mathcal{A}_n$ (resp. $\mathcal{A}'_E = \bigcup_{n=1}^{\infty} \mathcal{A}'_n$)—see Proposition A.1. Then the following result applies to this construction:

Proposition A.2. $E \sim^C E' \Rightarrow \mu_E(E) = \mu'_{E'}(E')$.

Proof. This is trivially true if $E \approx \widehat{\mathcal{A}}$. Assume, therefore that $\widehat{\mathcal{A}} \succ^C E, E'$. Let $N_{\mathrm{R},n}$ and $N_{\mathrm{NR},n}$ be the number of remainder and non-remainder events, respectively, in $\mathcal{D}_n$. If, for some n, there are no remainder events, then $N_{\mathrm{R},n} = 0$ and $N_{\mathrm{NR},n}$ is equal to the cardinality of $\mathcal{D}_n$. Define $N'_{\mathrm{R},n}$ and $N'_{\mathrm{NR},n}$ similarly for $\mathcal{D}'_n$. If $E_{\mathrm{R},n}, E_{\mathrm{NR},n} \in \mathcal{D}_n$ are remainder and non-remainder events, respectively, then $\mu_E(E_{\mathrm{NR},n}) = \frac{1-N_{\mathrm{R},n}\mu_E(E_{\mathrm{R},n})}{N_{\mathrm{NR},n}}$. Since a remainder event in $\mathcal{D}_n$ is a non-remainder event in $\mathcal{D}_{n+1}$, one has $\mu_E(E_{\mathrm{NR},n}) = \frac{1-N_{\mathrm{R},n}\mu_E(E_{\mathrm{NR},n+1})}{N_{\mathrm{NR},n}}$. This expression can be recursively substituted to give

$$\mu_E(E) = \frac{1}{N_{\mathrm{NR},0}} - \frac{N_{\mathrm{R},0}}{N_{\mathrm{NR},0}N_{\mathrm{NR},1}} + \frac{N_{\mathrm{R},0}N_{\mathrm{R},1}}{N_{\mathrm{NR},0}N_{\mathrm{NR},1}N_{\mathrm{NR},2}} - \frac{N_{\mathrm{R},0}N_{\mathrm{R},1}N_{\mathrm{R},2}}{N_{\mathrm{NR},0}N_{\mathrm{NR},1}N_{\mathrm{NR},2}N_{\mathrm{NR},3}} + \cdots . \tag{3}$$

Let $N_{\mathrm{NR},-1} \equiv 1$ and N_{n+1} be the number of non-remainder events that result when $E_{\mathrm{NR},n}$ is divided by $E_{\mathrm{R},n}$. Then

$$N_{\mathrm{R},n+1} = \begin{cases} 0 & \text{if there are no remainder events,} \\ N_{\mathrm{NR},n} & \text{otherwise,} \end{cases} \tag{4}$$

$$N_{\mathrm{NR},n+1} = N_{\mathrm{R},n} + N_{n+1}N_{\mathrm{NR},n} = N_{\mathrm{NR},n-1} + N_{n+1}N_{\mathrm{NR},n}. \tag{5}$$

Using $N_{\mathrm{R},0} = 1$, Eq. (3) becomes

$$\mu_E(E) = \frac{1}{N_{\mathrm{NR},0}} - \frac{1}{N_{\mathrm{NR},0}N_{\mathrm{NR},1}} + \frac{1}{N_{\mathrm{NR},1}N_{\mathrm{NR},2}} - \frac{1}{N_{\mathrm{NR},2}N_{\mathrm{NR},3}} + \frac{1}{N_{\mathrm{NR},3}N_{\mathrm{NR},4}} - \cdots,$$

and the series is finite and contains at most $n+1$ terms if and only if $N_{\mathrm{R},n} = 0$. The absolute value of the ratio of the $(n+1)$th to the nth terms is

$$\left| \frac{N_{\mathrm{NR},n}}{N_{\mathrm{NR},n+2}} \right| = \frac{1}{1 + N_{n+2}\frac{N_{\mathrm{NR},n+1}}{N_{\mathrm{NR},n}}} \leqslant \frac{1}{2}.$$

Thus the series expression for $\mu_E(E)$ converges and depends only on $\{N_{\mathrm{NR},n}\}$. By Lemma A.7 $\{N'_{\mathrm{NR},n}\} = \{N_{\mathrm{NR},n}\}$, thus $\mu_E(E) = \mu'_{E'}(E')$. □

Proposition A.3. *There exists a unique, solvable, and finitely additive measure representing* $\succcurlyeq^C$ *on* $\mathcal{A}$. *The measure* μ *is either atomless or purely and uniformly atomic.*

Proof. The arguments are trivial if $\mathcal{A}$ is finite. Thus assume otherwise, implying that $\mathcal{A}$ has no atoms. Consider any $E, E' \in \mathcal{A}$. If E is null, define $\mu(E) \equiv 0$. If E is not null, then define $\mu(E) = \mu_E(E)$, where μ_E is the measure in the construction preceding Proposition A.2. Because Proposition A.2 is true when $E = E'$, $\mu(E)$ is well defined (i.e., it does not depend on how the refinement with respect to E is constructed). Proposition A.2 establishes that $E \sim^C E' \Rightarrow \mu_E(E) = \mu'_{E'}(E')$. Suppose now that $\mu_E(E) = \mu_{E'}(E')$. If it is not the case that $E \sim^C E'$, then without loss of generality one can assume that $E \succ^C E'$. There is therefore some $e \in \mathcal{A}$ such that $e \subset E$, $E \succ^C e$, and $e \sim^C E'$. Since $E \backslash e$ is not null, it must be that $\mu_E(E) > \mu_e(e)$. To see this, note that $\mu_E(E)$ measures the number of times E 'fits' into $\widehat{\mathcal{A}}$, while $\mu_e(e)$ measures the number of times e 'fits' into $\widehat{\mathcal{A}}$, via the construction used in Proposition A.2. Using a similar construction, one can always find some $\mathcal{A}_E$ and $\mathcal{A}_e$, and $n > 0$ such that: e is strictly contained in both $a \in \mathcal{A}_{E,n}$ and $b \in \mathcal{A}_{e,n}$, both a and b are strictly contained in E, and $E \backslash a$ and $E \backslash b$ are not null. Thus, in both constructions, e fits more times into $\widehat{\mathcal{A}}$ than E, and it must be that $\mu_E(E) > \mu_e(e)$. From Proposition A.2 , $\mu_{E'}(E') = \mu_e(e) < \mu_E(E)$, which is a contradiction. Thus $\mu_E(E) = \mu'_{E'}(E') \Rightarrow E \sim^C E'$.

Assume now that $E \succ^C E'$. Then following an identical argument to the one in the previous paragraph one concludes that $\mu(E) > \mu(E')$. Now, if $\mu(E) > \mu(E')$ then it must be that $E \succ^C E'$, otherwise $E' \succcurlyeq^C E$ implying that $\mu(E) \leqslant \mu(E')$. Thus $\mu_E(E) = \mu'_{E'}(E') \Leftrightarrow E \sim^C E'$ and $\mu_E(E) > \mu'_{E'}(E') \Leftrightarrow E \succ^C E'$.

It remains to show that μ is additive: $\mu(E \cup E') = \mu(E) + \mu(E')$ when $E \cap E' = \emptyset$. If $E' \in \mathcal{A}_E$ for some algebra $\mathcal{A}_E$ constructed as in Propositions A.1 and A.2, then this immediately follows from the additivity of μ_E. Suppose otherwise, so that at least one of $\mathcal{A}_E$ and $\mathcal{A}_{E'}$ must be infinite and fine. Assume, without loss of generality that $\mathcal{A}_E$ is infinite, and that for some disjoint $E, E' \in \mathcal{A}$, $\mu(E \cup E') > \mu(E) + \mu(E')$. Then one can find some $\mathcal{A}_E$ and $n > 0$ such that: E' is strictly contained in $a \in \mathcal{A}_{E,n}$, with $E \cap a = \emptyset$, $a \backslash E'$ not null, and $\mu(E \cup a) = \mu(E) + \mu(a)$ (by additivity of μ_E). Appealing to the same argument used earlier, $\mu(a) > \mu(E')$ and $\mu(E \cup E') < \mu(E \cup a)$. Because $\mathcal{A}_E$ contains no atoms, μ_E is fine and tight (see [4]) and one can choose a so that $\mu(E \cup E') > \mu(E) + \mu(a) = \mu(E \cup a)$, leading to a contradiction. A similar argument rules out $\mu(E \cup E') < \mu(E) + \mu(E')$.

The solvability of μ follows from the fact that $\mathcal{A}$ is homogeneous. □

The fact that $\mathcal{A}$ is homogeneous implies that f and f' induce the same lottery with respect to μ on $\mathcal{A}$, then $f \widehat{\mathcal{A}} g \sim f' \widehat{\mathcal{A}} g$ for every $g \in \mathcal{F}$. □

Proof of Proposition 1. See the proof to Proposition 2 in Sagi [32]. □

Proof of Theorem 2. Theorem 1 implies that the decision maker is indifferent between any $f \widehat{\mathcal{A}} h$ and $g \widehat{\mathcal{A}} h$ whenever $L_{\mu,f} = L_{\mu,g}$. It is therefore sufficient to establish that, for each h, $\succcurlyeq$ induces a continuous representation over $\overline{\Delta(\mu)}$. The proof proceeds in four steps. First, it is shown that every weakly convergent sequence in $\Delta(\mu)$, say $\{L_n\} \to L$, can be associated with a sequence of acts, $\{f_n\}$, f, such that $L_{\mu,f} = L$, $L_{\mu,f_n} = L_n$, and $\{f_n\}$ μ-converges in probability to f. I.e., for each $\varepsilon > 0$, $\lim_{n\to\infty} \mu\Big(\big\{\omega|\varepsilon < d\big(f_n(\omega), f(\omega)\big)\big\}\Big) = 0$. Next, it is shown that μ is countably additive. In the third step, the previous two steps are used to establish that $V(f_n) \to V(f)$, where $V(\cdot)$ is the continuous utility function in Proposition 1. Finally, this is used to construct a continuous function, U_h, representing $\succcurlyeq$ on $\overline{\Delta(\mu)}$.

Step 1: Begin by arbitrarily selecting some $f \in \mathcal{F}$ such that $L_{\mu,f} = L$, and everywhere outside of $\widehat{\mathcal{A}}$, $f = h$ (by definition, there must be such an act). Let $X_f = \{x_1, \ldots, x_N\}$ correspond to the finite range of $f(\widehat{\mathcal{A}})$, and let $E_i = f^{-1}(x_i) \cap \widehat{\mathcal{A}}$ and $p_i = \mu(E_i)$, for $i = 1, \ldots, N$. Denoting the metric on X as $d(\cdot,\cdot)$, there is some $\delta > 0$ such that $d(x_i, x_j) > \delta$ for all distinct $x_i, x_j \in X_f$. For each n, let $Y_n \subset X$ be the set of payoffs of L_n. Let $q_{j,n}$ be the probability assigned to $y_j \in Y_n$ by L_n, and for each $i = 1, \ldots, N$, let

$$p_{i,n} = \sum_{y_j \in Y_n, d(x_i, y_j) < \delta} q_{j,n},$$

while $p_{A,n} = 1 - \sum_{i=1}^{N} p_{i,n}$.

Construct f_n as follows: If $p_i \geqslant p_{i,n}$, for any $i = 1, \ldots, N$, then f_n assigns each $y_j \in Y_n$ with $d(y_j, x_i) < \delta$ to some arbitrary $E_{ij,n} \subseteq E_i$ with $\mu(E_{ij,n}) = q_{j,n}$ (solvability of μ implies one can always do this). Let $a_{i,n} = E_i \setminus \bigcup_j E_{ij,n}$. If $p_i < p_{i,n}$ then f_n assigns outcomes $y_j \in Y_n$ with $d(y_j, x_i) < \delta$ to subevents of $E_i \cup b_{i,n}$ where $b_{i,n} \subseteq \bigcup_i a_{i,n}$, and $\mu(E_i \cup b_{i,n}) = p_{i,n}$ ($b_{i,n}$ is otherwise arbitrary). The remaining outcomes in Y_n are assigned to corresponding disjoint events in $\widehat{\mathcal{A}}$. Everywhere outside of $\widehat{\mathcal{A}}$ define $f_n = f = h$.

Weak convergence of L_n to L implies that $p_{i,n} \to p_i$ for each $i = 1, \ldots, N$, while $p_{A,n} \to 0$. If $\{f_n\}$ does not μ-converges to f in probability, then the former limits imply that there is some $x_i \in X_f$, $\varepsilon > 0$, and a subsequence of $\{f_n\}$, say $\{g_n\}$, such that $\mu\Big(\big\{\omega|\varepsilon < d\big(g_n(\omega), x_i\big) < \delta\big\}\Big) > \eta > 0$ for all n. The latter is not consistent with the weak convergence of L_{μ,g_n} to L.

Step 2: To prove μ is countably additive, it is sufficient to show that for any sequence $\{E_n\} \subset \Sigma$ such that $E_{n+1} \subseteq E_n$ and $\cap_n E_n = \emptyset$, $\mu(E_n) \to 0$ (see, for example, Theorem 1.2.8 from Ash [2]). If $\mu(E_n) \not\to 0$, then there is some subsequence of decreasing events, $\{E'_n\} \subseteq \{E_n\}$ such that $\mu(E'_n) \geqslant \varepsilon$, for some $\varepsilon > 0$ and all n. Assume, for now, that $\mu(E'_n) < 1$ for some finite j. Then because μ has no atoms, one can always find some non-null $A \in \mathcal{A}$ such that $0 < \mu(A) < \varepsilon$ and $A \cap E'_n = \emptyset$ for all $n > j$. Homogeneity of $\mathcal{A}$ implies that each E'_n, where $n > j$, contains some $e_n \in \mathcal{A}$, so that $xe_ny \sim xAy$ for all $x, y \in X$. Thus $V(xe_ny) \to V(xAy)$ for every $x, y \in X$. On the other hand, $\bigcap_n e_n = \emptyset$, implying that for every $\omega \in S$, there must be some k such that $\omega \notin E'_k$ and therefore not in any e_{k+m} for all $m \geqslant 0$. In turn, it must be that xe_ny converges pointwise to y, thus by Proposition 1, $V(xe_ny) \to V(y)$, which because A is not null, contradicts Axiom N′.

It remains to examine the case $\mu(E'_n) = 1$ for all n. In this case, consider any $A \in \mathcal{A}$ such that $\mu(A) < 1$ and construct the sequence $\{E'_n \cap A\}$. One can apply the previous argument to this sequence, noting that $\mu(E'_n \cap A) = \mu(A)$ for all n, and $\bigcap_n (E'_n \cap A) = \emptyset$.

Step 3: A minor extension of Theorem 2.5.3 in Ash [2] (its proof applies more generally to separable metric spaces) implies that every infinite subsequence of $\{f_n\}$ contains another subsequence, say $\{\hat{f}_n\}$, that converges μ-a.e. to f. In turn, this implies that $V(\hat{f}_n)$ from Proposition 1 converges to $V(f)$ (because X is compact, and points of non-convergence form a null set of μ and therefore of $\succcurlyeq$). If $V(f_n)$ does not converge to $V(f)$, then there is a subsequence of $\{f_n\}$, say $\{g_n\}$, such that $|V(g_n) - V(f)| > \eta > 0$ for all n. The subsequence, $\{g_n\}$, however, must contain another subsequence, $\{\hat{f}_n\}$ for which V does converge—a contradiction. Thus $V(f_n) \to V(f)$.

Step 4: Fix $h \in \mathcal{F}$. For each $L \in \Delta(\mu)$, find some $f \in \mathcal{F}$, such that $L = L_{\mu,f}$ (this can always be done by definition of $\Delta(\mu)$). set $U_h(L) = V(f \overset{\frown}{\mathcal{A}} h)$. Then U_h represents $\succcurlyeq$ and, by steps 1-3, is continuous on $\Delta(\mu)$. Because $\mathcal{F}$ is dense in X^{Ω}, this continuous representation can be extended to $\overline{\Delta(\mu)}$. □

References

[1] M. Allais, Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'cole amricaine, Econometrica 21 (1953) 503–546.
[2] R. Ash, Real Analysis and Probability, Academic Press, New York, 1972.
[3] S.H. Chew, A generalization of the quasilinear mean with applications to the measurement of income inequality and decision theory resolving the Allais paradox, Econometrica 51 (1983) 1065–1092.
[4] S.H. Chew, J.S. Sagi, Event exchangeability: probabilistic sophistication without continuity or monotonicity, Econometrica 74 (2006) 771–786.
[5] B. de Finetti, Foresight: its logical laws, its subjective sources, Ann. Inst. Henri Poincaré (1937) and translated in: H.E. Kyburg, H.E. Smokler (Eds.), Studies in Subjective Probability, Wiley, New York, 1964.
[6] D. Ellsberg, Risk, ambiguity and the Savage axioms, Quart. J. Econ. 75 (1961) 643–669.
[7] L.G. Epstein, A definition of uncertainty aversion, Rev. Econ. Stud. 66 (1999) 579–608.
[8] L.G. Epstein, J. Zhang, Subjective probabilities on subjectively unambiguous events, Econometrica 69 (2001) 265–306.
[9] H. Ergin, F. Gul, A subjective theory of compound lotteries, mimeo, 2004.
[10] W. Fellner, Distortion of subjective probabilities as a reaction to uncertainty, Quart. J. Econ. 75 (1961) 670–689.
[11] C. Fox, A. Tversky, Ambiguity aversion and comparative ignorance, Quart. J. Econ. 199 (1995) 585–603.
[12] P. Ghirardato, M. Marinacci, Ambiguity made precise: a comparative foundation, J. Econ. Theory 102 (2002) 251–289.
[13] P. Ghirardato, M. Marinacci, F. Maccheroni, Differentiating ambiguity and ambiguity attitude, J. Econ. Theory 118 (2004) 133–173.
[14] I. Gilboa, D. Schmeidler, Maximin expected utility with a non-unique prior, J. Math. Econ. 18 (1989) 141–153.
[15] S. Grant, Subjective probability without monotonicity: or how machina's mom may also be probabilistically sophisticated, Econometrica 63 (1995) 159–191.
[16] J. Gyntelberg, F. Hansen, Expected utility theory with "small worlds", mimeo, 2004.
[17] C. Heath, A. Tversky, Preference and belief: ambiguity and competence in choice under uncertainty, J. Risk Uncertainty 4 (1991) 5–28.
[18] M. Hsu, M. Bhatt, R. Adolphs, D. Tranel, C.F. Camerer, Neural systems responding to degrees of uncertainty in human decision-making, Science 310 (2005) 1680–1683.
[19] H.-J. Keppe, M. Weber, Judged knowledge and ambiguity aversion, Theory Dec. 39 (1995) 51–77.
[20] J.M. Keynes, A Treatise on Probability, Macmillan, London, 1921.
[21] P.M. Klibanoff, M. Marinacci, S. Mukerji, A smooth model of decision making under uncertainty, Econometrica 73 (2005) 1849–1892.
[22] F.H. Knight, Risk, Uncertainty and Profit, Houghton Mifflin, Boston, 1921.
[23] I. Kopylov, Subjective probabilities on "small" domains, J. Econ. Theory 133 (2007) 236–265.
[24] D.M. Kreps, Notes on the Theory of Choice, Westview Press, London-Boulder, 1988.
[25] M.J. Machina, D. Schmeidler, A more robust definition of subjective probability, Econometrica 60 (1992) 745–780.
[26] Y. Nakamura, Subjective expected utility with non-additive probabilities on finite state spaces, J. Econ. Theory 51 (1990) 346–366.
[27] R. Nau, Uncertainty aversion with second-order probabilities and utilities, Manage. Sci. 52 (2006) 136–145.
[28] K. Nehring, Decision-making in the context of imprecise probabilistic beliefs, mimeo, 2006.
[29] K. Nehring, Bernoulli without bayes: a theory of utility-sophisticated preferences under ambiguity, mimeo, 2006.
[30] K. Nehring, Is it possible to define subjective probabilities in purely behavioral terms? A comment on Epstein–Zhang (2001), mimeo, 2006.
[31] F.P. Ramsey, Truth and probability, in: R.B. Braithwaite (Ed.), The Foundations of Mathematics and Other Logical Essays, Routledge, London, 1931(Chapter 7).
[32] J. Sagi, Modeling implications of source-invariance to Machina's 'almost objective fair bets', mimeo, 2007.
[33] J.L. Savage, The Foundations of Statistics, Wiley, New York, 1954.
[34] D. Schmeidler, Subjective probability and expected utility without additivity, Econometrica 57 (1989) 571–587.

[35] M. Schroder, C. Skiadas, Optimal lifetime consumption-portfolio strategies under trading constraints and non-tradeable income, Stochastic Processes Appl. 108 (2003) 155–202.
[36] M. Siniscalchi, A behavioral characterization of plausible priors, J. Econ. Theory 128 (2006) 91–135.
[37] C. Skiadas, Dynamics portfolio choice and risk aversion, in: J.R. Birge, V. Linetsky (Eds.), Financial Engineering, vol. 15, Elsevier, Amsterdam, 2008.
[38] A. Tversky, P.P. Wakker, Risk attitudes and decision weights, Econometrica 63 (1995) 1255–1280.
[39] P.P. Wakker, Uncertainty, in: The New Palgrave: A Dictionary of Economics, Palgrave-Macmillan, 2007, forthcoming.
[40] J. Zhang, Qualitative probabilities on λ-systems, Math. Soc. Sci. 38 (1999) 11–20.