Vector-Adjusted Expected Utility

Marciano Siniscalchi*
Princeton University

First Draft: May 2000
This Draft: February 2001

Abstract

This paper proposes a representation of (possibly) probabilistically unsophisticated preferences whereby (1) beliefs are jointly represented by a finitely additive probability measure and a vector-valued measure; (2) uncertain prospects are ranked according to the difference between a baseline expected utility evaluation and an adjustment term; and (3) the latter is the norm of the vector-valued expected utility of the prospect under consideration.

Vector-valued measures are employed to represent the extent to which ambiguity about different events "cancels out" or "adds up", as revealed by the decision maker's preferences.

The proposed representation, *vector-adjusted expected utility* (VEU), is shown to be consistent with the maxmin-expected utility model (MEU). A necessary and sufficient condition characterizing the class of VEU preferences within the MEU family of preferences is provided.

1 Introduction

According to the Bayesian paradigm of choice under uncertainty, individuals express their subjective beliefs about the relative likelihood of events by assigning probabilities to them. Correspondingly, they rank uncertain prospects according to their subjective expected utility.

Daniel Ellsberg [5] first observed that, in some circumstances, the relative likelihood of events may be perceived as being "ambiguous". He also noted that a pessimistic (or optimistic) attitude in the face of ambiguity may lead to patterns of preferences that do not fit within the Bayesian

*Department of Economics, Princeton University, Princeton, NJ, 08544-1021. Email: marciano@princeton.edu.
An early draft of this paper was titled "A Vector-Valued Representation of Knightian Uncertainty." I would like to thank Paolo Ghirardato, Faruk Gul, Alessandro Lizzeri, Fabio Maccheroni, Massimo Marinacci, and Josè Scheinkman for useful discussion. All errors are my own. Parts of this paper were written while visiting IGIER - Università Bocconi and the MEDS Department, Kellogg Graduate School of Management, Northwestern University.

paradigm. In order to account for such deviations, generalizations of the Bayesian model of choice have been proposed and axiomatized: most notably, the Choquet-Expected Utility (Schmeidler [18]) and Maxmin-Expected Utility (Gilboa and Schmeidler [11]) models.

This paper suggests an alternative extension of the Bayesian model, *vector-adjusted expected utility*. In the proposed representation, the decision maker's subjective beliefs are (uniquely) separated into a baseline probabilistic component and an adjustment term, which reflects the perceived distortions (e.g. ambiguity) that motivate deviations from the Bayesian model. Correspondingly, an "ambiguity-averse" decision maker will rank uncertain prospects according to the difference between an expected utility component and an adjustment component. I provide a characterization of the proposed decision rule, along with comparisons with other models of choice under uncertainty.

The essential features of vector-adjusted expected utility (VEU henceforth) are best illustrated in the context of one of Ellsberg's well-known "mind experiments." A decision maker is presented with an urn containing a total of 90 balls; of these, 30 are red, and the remaining are green or blue, in unspecified proportions. The decision maker is asked to rank the following prospects, or "bets", labelled $r, g, \bar{r}$ and $\bar{g}$:

(r) receive \$10 if a red ball is drawn and \$0 otherwise;

(g) receive \$10 if a green ball is drawn, and \$0 otherwise;

($\bar{r}$) receive \$10 if a red ball is *not* drawn, and \$0 otherwise;

($\bar{g}$) receive \$10 if a green ball is *not* drawn, and \$0 otherwise.

Several individuals, when faced with these options, prefer r to g and $\bar{r}$ to $\bar{g}$. Recall that the proportion of green vs. blue balls is unspecified; an "ambiguity averse" individual may then prefer to bet on a red draw rather than on a green draw, because she is worried that the urn may contain no green balls. By a symmetric argument, she may simultaneously prefer to bet *against* a red draw, rather than against a green draw, because she is worried that the urn may contain 60 green balls.

One can verify that, if a decision maker expresses these preferences, then her beliefs cannot be represented by a probability measure (the choice of utility index is immaterial). However, these preferences (and beliefs) admit a VEU representation.

To see this, define a *baseline probability measure* p_1 by letting $p_1(\rho) = p_1(\gamma) = p_1(\beta) = \frac{1}{3}$, where ρ, γ and β denote the elementary events whereby a red, green or blue ball is drawn.

Next, in order to represent the adjustments to the probabilistic assessment p_1, define a *signed measure* p_2 on $\{\rho, \gamma, \beta\}$ by letting $p_2(\rho) = 0$, $p_2(\gamma) = \frac{1}{3}$ and $p_2(\beta) = -\frac{1}{3}$, and extending it to all subsets of $\{\rho, \gamma, \beta\}$ by additivity. Thus, for instance, $p_2(\rho \cup \beta) = -\frac{1}{3}$.

The signed measure p_2 is to be viewed as a vector-valued measure (albeit a particularly simple one), whose range is the vector space of the reals endowed with the absolute value norm. The general VEU representation allows for arbitrary vector-valued measures.

Finally, fix a utility function $u(\cdot)$ on the relevant space of outcomes, such that $u(\$10) > u(\$0)$. It can then be verified that the VEU functional associating to each prospect f the quantity

$$\sum_{\omega \in \{\rho,\gamma,\beta\}} u(f(\omega))p_1(\omega) - \Big| \sum_{\omega \in \{\rho,\gamma,\beta\}} u(f(\omega))p_2(\omega) \Big|.$$

represents the preferences indicated above.

The interpretation of the adjustment p_2 for an ambiguity-averse decision maker is as follows. First, for any event $E \subset \{\rho, \gamma, \beta\}$, the *absolute value* of $p_2(E)$ reflects, loosely speaking, the extent to which ambiguity reduces the decision maker's willingness to bet on E. For example, the adjustments $p_2(\rho) = 0$ and $p_2(\gamma) = \frac{1}{3}$ signal that there is no (perceived) ambiguity about ρ, but there is substantial (perceived) ambiguity about γ; this, in turn, makes a bet on γ less palatable than a bet on ρ, although the two events are assigned the same baseline probability.

Second, for any event $E \subset \{\rho, \gamma, \beta\}$, the *sign* of $p_2(E)$ indicates how ambiguity about E combines with ambiguity about other, disjoint events. For instance, $p_2(\gamma)$ and $p_2(\beta)$ are equal in absolute value, but have opposite signs; this indicates that the events γ and β are perceived as being equally ambiguous, but their ambiguities exactly "cancel out": $p_2(\gamma \cup \beta) = 0$. On the other hand, ambiguity about β is preserved when taking its union with ρ: $p_2(\rho \cup \beta) = -\frac{1}{3}$.

The example emphasizes the two main features of the VEU model of choice under uncertainty. First, it portrays a decision maker who evaluates a prospect by formulating a "neutral" assessment of the latter, and separately accounting for any perceived ambiguity. Second, it provides an explicit, compact representation of the extent to which ambiguity about disjoint events "cancels out" or "adds up" when taking unions. As the example illustrates, these *patterns of interaction* are closely related to the deviations from the Bayesian paradigm one observes as a consequence of (perceived) ambiguity. I suggest that both features may contribute to making VEU an appealing modelling option in the analysis of other, more complex problems of choice under uncertainty.

Observe also that the general VEU representation involves a linear functional, an integral with respect to a vector-valued measure, and a well-behaved convex function of the latter. In sufficiently regular setting, vector integrals can be thought of simply as collections of scalar integrals; hence, they are structurally very similar to standard expectation operators. As a result, VEU generally provides an analytically and numerically tractable representation of preferences.

The characterization of VEU preferences provided in this paper takes place in the setting of Anscombe and Aumann [1]. I first show that every VEU preference also admits a Maxmin-Expected Utility (MEU) representation: that is, preferences are also represented by the functional associating to each act its minimum expected utility taken over a closed, convex set of probability measures. The main result of this paper then establishes that a single additional axiom characterizes the VEU subclass of preferences within the MEU class.

The paper is organized as follows. Section 2 defines vector measures and integrals, and provides the necessary decision-theoretic background. Section 3 contains the main characterization result. Section 4 provides an extended discussion and interpretation of the axioms; moreover, it contains additional results relating VEU preferences to other models of choice, as well as further examples. All proofs are in the Appendix.

2 Preliminaries and Background

2.1 Vector Measures and Integration

Consider a set Ω, an algebra Σ of subsets of Ω, and a real Banach space $\mathcal{V}$ with norm $\|\cdot\|_{\mathcal{V}}$; denote by $0_{\mathcal{V}}$ the zero vector in $\mathcal{V}$. A *vector measure* (cf. [4], §III.1.1) is a finitely additive set function $\mu : \Sigma \to \mathcal{V}$; that is, for disjoint $E, F \in \Sigma$, $\mu(E \cup F) = \mu(E) + \mu(F)$, where "+" indicates addition in $\mathcal{V}$. The set of $\mathcal{V}$-valued vector measures on (Ω, Σ) will be denoted by $\Delta^{\mathcal{V}}(\Omega, \Sigma)$; the set of finitely additive probability measures on the same measurable space is denoted by $\Delta(\Omega, \Sigma)$.

For the purposes of the present paper, it is sufficient to define integration of bounded, real-valued functions with respect to a finitely additive vector-valued measure. The material presented here is based on Bartle [3]; the interested reader is referred to the latter paper for further details and a considerably more general treatment.

Denote by $B_0(\Omega, \Sigma)$ the set of simple Σ-measurable real-valued functions on Ω (i.e. the set of finite linear combinations of indicator functions), and let $B(\Omega, \Sigma)$ denote its uniform closure. As is customary, for any $c \in \mathbb{R}$, let "c" also denote the function assigning the value c to each $\omega \in \Omega$. The indicator function of an event $E \in \Sigma$ will be denoted by 1_E. Also, to simplify notation, the sets $B_0(\Omega, \Sigma)$ and $B(\Omega, \Sigma)$ will be simply referred to as B_0 and B, respectively, unless it is necessary to emphasize the dependence on Ω and/or Σ.

For any simple function $f \in B_0$, the $\mathcal{V}$-valued integral with respect to a $\mathcal{V}$-valued measure μ is defined as in in the development of Lebesgue integration: if $f = \sum_{k=1}^{K} x_k 1_{E_k}$ for some collection of disjoint events $E_1, \ldots, E_K$ and nonzero real numbers $x_1, \ldots, x_K$, then, for any event $E \in \Sigma$,

$$I(f; \mu, E) = \int_E f \mu(d\omega) = \sum_{k=1}^{K} \mu(E_k \cap E) x_k. \tag{1}$$

It is routine to show that the above definition is independent of the representation of the integrand, and that the integral thus defined is a linear operator on B_0 with values in $\mathcal{V}$, and an additive function on Σ for any fixed $f \in B_0$.

For non-simple functions $f \in B \setminus B_0$, by definition there exists a sequence $\{f_n\}_{n \geq 1} \subset B_0$ converging uniformly to f; then the integral of f w.r.to a $\mathcal{V}$-valued measure μ on an event $E \in \Sigma$

is defined by

$$I(f;\mu,E)=\int_E f\mu(d\omega)=\lim_{n\to\infty}\int_E f_n\mu(d\omega), \tag{2}$$

where the above limit is taken with respect to the norm topology on $\mathcal{V}$; see Bartle [3], Theorem 8.

Example 1 Let $\mu=(\mu_1,\ldots,\mu_n)$ be a collection of bounded, finitely additive signed scalar measures on (Ω,Σ), and let $\mathcal{V}=\mathbb{R}^n$, endowed with the norm $\|x\|=(\sum_n |x_n|^p)^{1/p}$, $1\le p\le\infty$. Then μ is an l^p_n-valued vector measure.

In this case, $\int_E f\mu(d\omega)=(\int_E f\mu_i(d\omega))_{i=1}^n\in\mathbb{R}^n$, where each component of the latter n-vector is a Dunford integral (cf. [4], §III.2).

The VEU representation of preferences involves a finitely additive probability measure, as well as a vector measure. The two measures must satisify a mutual consistency requirement; also, a normalization condition must be imposed on the vector measure. The following definition provides the details.

Definition 1 *Consider a finitely additive probability measure $p_1\in\Delta(\Omega,\Sigma)$ and a vector measure $p_{-1}\in\Delta^{\mathcal{V}}(\Omega,\Sigma)$. Then p_{-1} is a* vector adjustment for p_1 *if and only if:*
(i) For all $E\in\Sigma$, $p_1(E)\ge\|p_{-1}(E)\|_{\mathcal{V}}$;
(ii) $p_{-1}(\emptyset)=p_{-1}(\Omega)=0_{\mathcal{V}}$.

Note that the definition implies that p_{-1} is bounded.[1]

Since $\|p_{-1}(E)\|_{\mathcal{V}}$ is interpreted as an "adjustment" to the probabilistic assessment $p_1(E)$, Condition (i) ensures that this adjustment results in a non-negative decision weight. Condition (ii) reflects the obvious requirement that there be no ambiguity regarding the decision weights assigned to the empty set or the entire space.

2.2 Decision-Theoretic Setup

Following Anscombe-Aumann [1], Schmeildler [18] and Gilboa-Schmeidler [11], consider a set of states of nature Ω endowed with an algebra Σ, a set X of consequences (prizes), the set Y of (finite-support) lotteries on X, and a convex subset $L\subset Y^\Omega$ of Σ-measurable acts, including the class L_0 of simple functions (i.e. functions taking up finitely many values in Y), which in turn

[1] Condition (i) of Definition 1 (along with the normalization condition $p_1(\Omega)=1$) also ensures that the total variation (cf. Dunford-Schwartz [4], §III.1.4), and hence the semi-variation ([4], §IV.10.3) of the vector-valued measure p_{-1} are finite. This simplifies the treatment of integration of *un*bounded functions with respect to vector-valued measures: see e.g. the discussion of Definition 1 in [3].

includes the class L_c of all constant functions. With the usual abuse of notation, denote by y the constant act assigning the lottery $y \in Y$ to each $\omega \in \Omega$.

Next, consider a preference relation $\succeq$ defined on L_c, and denote by $L(\succeq)$ the subset of Y^Ω comprising all acts that are bounded for $\succeq$: that is, an act $f \in Y^\Omega$ is an element of $L(\succeq)$ if and only if there exist lotteries $y, y' \in Y$ such that $y \succeq f \succeq y'$.

Mixtures of acts are taken pointwise: for every pair of acts $f, g \in L$ and for any $\alpha \in [0,1]$, $\alpha f + (1-\alpha)g$ is the act assigning the compound lottery $\alpha f(\omega) + (1-\alpha)g(\omega)$ to each state $\omega \in \Omega$.

Gilboa and Schmeidler [11] consider the following axioms:

(i) WEAK ORDER. $\succeq$ is transitive and complete.

(ii) MONOTONICITY. $\forall f, g \in L$, $f(\omega) \succeq g(\omega)$ implies $f \succeq g$.

(iii) CONTINUITY. $\forall f, g, h \in L$ such that $f \succ g \succ h$, there exist $\alpha, \beta \in (0,1)$ such that $f \succ \alpha f + (1-\alpha)h \succ g$ and $g \succ \beta f + (1-\beta)h \succ h$.

(iv) CONSTANT-ACT INDEPENDENCE. $\forall f, g \in L$, $h \in L_c$ and $\alpha \in (0,1)$: $f \succeq g$ implies $\alpha f + (1-\alpha)h \succeq \alpha g + (1-\alpha)h$.

(v) AMBIGUITY AVERSION. $\forall f, g \in L$ and $\alpha \in (0,1)$: $f \sim g$ implies $\alpha f + (1-\alpha)g \succeq g$.

(vi) NON-DEGENERACY. Not for all $f, g \in L$, $f \succeq g$.

The following result characterizes the MEU decision rule in the setting under consideration (see Klibanoff *et al.* [14] for a similar result in a setting without objective lotteries).

Theorem 2.1 (Gilboa and Schmeidler [11]) *For any preference relation $\succeq$ on L_0, the following statements are equivalent:*

(1) $\succeq$ satisfies Axioms (i)-(v) for $L = L_0$;

(2) there exists a convex and weak-closed set of priors $\mathcal{C} \subset \Delta(\Omega, \Sigma)$ and an affine utility function $u : Y \to \mathbb{R}$ (unique up to a p.l.t.) such that, for all acts $f, g \in L_0$,*

$$f \succeq g \quad \Leftrightarrow \quad \min_{q \in \mathcal{C}} \int_\Omega u \circ f \, dq \geq \min_{q \in \mathcal{C}} \int_\Omega u \circ g \, dq. \tag{3}$$

Moreover, if (1) above holds, then $\succeq$ has a unique extension to $L(\succeq)$.

Finally, $\mathcal{C}$ is unique if and only if Axiom (vi) holds.

Note that a preference relation $\succeq$ admits a standard EU representation if and only if it satisfies Axioms (i)-(iv), plus

(v') AMBIGUITY NEUTRALITY. $\forall f, g \in L$ and $\alpha \in (0,1)$: $f \sim g$ implies $\alpha f + (1-\alpha)g \sim g$.

The following simple results will play an important role in the next section.

Remark 2.1 (Solvability) *If a preference relation $\succeq$ satisfies Axioms (i), (iii) and (iv)[2], then, for*

[2] It is actually sufficient to require that $\succeq$ satisfy Independence on L_c.

every act $f \in L(\succeq)$, there exists a lottery $y(f) \in Y$ such that $y(f) \sim f$.

To simplify the statement of axioms and results, for any act $f \in L(\succeq)$ the notation $y(f)$ will always indicate a lottery such that $y(f) \sim f$; the preceding remark asserts the existence of at least one such lottery.

Remark 2.2 *If a preference relation $\succeq$ satisfies Axioms (i)-(iv), then it satisfies Axiom (v), Ambiguity Aversion (resp. Axiom (v'), Ambiguity Neutrality), if and only if, for every $f, g \in L$ and $\gamma \in [0,1]$,*

$$\gamma f + (1-\gamma) g \ \succeq (\sim)\ \gamma y(f) + (1-\gamma) y(g).$$

Finally, the expression "VEU (resp. MEU, EU) preference" indicates a preference relation that admits a VEU (resp. MEU, EU) representation.

3 Characterization of Vector-Adjusted Expected Utility

In the interest of clarity, the characterization of VEU preferences is presented here in two parts. Proposition 3.1 below provides the first: it states that every VEU preference relation also admits a MEU representation.

Proposition 3.1 *Consider a finitely additive probability measure $p_1 \in \Delta(\Omega, \Sigma)$, a Banach space $(\mathcal{V}, \|\cdot\|_{\mathcal{V}})$, a vector adjustment $p_{-1} \in \Delta^{\mathcal{V}}(\Omega, \Sigma)$ for p_1 and an affine utility function $u : Y \to \mathbb{R}$. If, for a preference relation $\succeq$ on L, and for all acts $f, g \in L(\succeq)$,*

$$f \succeq g \quad \Leftrightarrow \quad \int_\Omega u \circ f \; p_1(d\omega) - \Big\| \int_\Omega u \circ f \; p_{-1}(d\omega) \Big\|_{\mathcal{V}} \geq \int_\Omega u \circ g \; p_1(d\omega) - \Big\| \int_\Omega u \circ g \; p_{-1}(d\omega) \Big\|_{\mathcal{V}} \quad (4)$$

then there exists a weak-closed and convex set $\mathcal{C} \subset \Delta(\Omega, \Sigma)$ such that Equation (3) in Theorem 2.1 holds for all acts $f, g \in L(\succeq)$.*

The second part of the characterization shows that a single additional axiom, *Symmetric Hedging*, identifies the subclass of VEU preferences within the MEU family of preferences.

The reader is referred to Section 4 for an extended discussion of this axiom and related definitions; its basic motivation may be summarized as follows.

Ambiguity Aversion postulates a (weak) preference for "'smoothing' or averaging utility distributions" (Schmeidler [18], p. 582). This preference, in turn, can be interpreted as reflecting a desire to "hedge against ambiguity": a decision maker who is unable to confidently assess the relative likelihood of certain events may wish to reduce the perceived variability of outcomes (i.e. the variability of utilities) contingent on them.

If one accepts the hedging motivation, then it follows that, whenever two pairs of acts provide the same opportunities for reducing perceived outcome variability across the same events, the decision maker's preferences should reveal that, indeed, mixing the two pairs of acts provides an equally effective "hedge against ambiguity." This is, essentially, what Symmetric Hedging requires.

In order to state the latter axiom, a preliminary definition is required.

Definition 2 *Two acts $f, \bar{f} \in L$ are* complementary *if and only if, for any two states $\omega, \omega' \in \Omega$,*

$$\frac{1}{2}f(\omega) + \frac{1}{2}\bar{f}(\omega) \sim \frac{1}{2}f(\omega') + \frac{1}{2}\bar{f}(\omega').$$

Thus, if an act f assigns different outcomes at two states, then so does a complementary act $\bar{f}$. Moreover, changes in the outcome assigned by f are exactly offset by the corresponding changes in the outcome assigned by $\bar{f}$. Note also that, if the pairs of acts $f, \bar{f}$ and, respectively, $g, \bar{g}$ are complementary, then, for any weight $\gamma \in [0, 1]$, the mixtures $\gamma f + (1-\gamma)g$ and $\gamma\bar{f} + (1-\gamma)\bar{g}$ are themselves complementary. This implies that f and g provide the same opportunities for utility smoothing as do $\bar{f}$ and $\bar{g}$. Section 4.1.1 provides further details.

The Symmetric Hedging axiom may now be formulated.[3]

Symmetric Hedging. For all acts $f, g, \bar{f}, \bar{g} \in L$ and real numbers $\gamma \in [0, 1]$: if the pairs $f, \bar{f}$ and, respectively, $g, \bar{g}$ are complementary, and there exist a real number $\alpha \in (0, 1)$ and lotteries $y, y' \in Y$ such that

$$\alpha\Big(\gamma f + (1-\gamma)g\Big) + (1-\alpha)y \sim \alpha\Big(\gamma y(f) + (1-\gamma)y(g)\Big) + (1-\alpha)y' \tag{5}$$

then

$$\alpha\Big(\gamma\bar{f} + (1-\gamma)\bar{g}\Big) + (1-\alpha)y \sim \alpha\Big(\gamma y(\bar{f}) + (1-\gamma)y(\bar{g})\Big) + (1-\alpha)y'. \tag{6}$$

Recall that, by Remark 2.2, a MEU decision maker will weakly prefer a mixture of acts to the same mixture of lotteries equivalent to them; as noted above, this captures a notion of hedging against ambiguity.

Equation 5 indicates that mixing $\gamma f + (1-\gamma)$ with y and $\gamma y(f) + (1-\gamma)y(g)$ with y' exactly offsets the decision maker's (weak) preference for $\gamma f + (1-\gamma)g$ vs. $\gamma y(f) + (1-\gamma)y(g)$; that is, it counterbalances the hedging effect of mixing f and g.

Equation 6 similarly indicates that the *same* lotteries also exactly offset the hedging effect of mixing $\bar{f}$ and $\bar{g}$. Thus, Symmetric Hedging encodes the requirement that f, g and $\bar{f}, \bar{g}$ provide an equally effective hedge against ambiguity.

Observe that, by Remark 2.2, EU preferences always satisfy Symmetric Hedging.

[3] I thank Paolo Ghirardato for suggesting the present formulation of the axiom.

Further remarks about the implications of Symmetric Hedging are offered in Section 4.1.2; I now turn to the main characterization result. Recall that, for any vector measure μ, $I(\cdot;\mu,\Omega)$ denotes the vector integral operator $f \mapsto \int_\Omega f\mu(d\omega)$, and $I(B;\mu;\Omega)$ denotes the image of B through $I(\cdot;\mu,\Omega)$.

Theorem 3.2 *Consider a preference relation $\succeq$ for which there exist an affine utility function $u : Y \to \mathbb{R}$ and a weak*-closed and convex set $\mathcal{C} \subset \Delta(\Omega,\Sigma)$ such that Equation (3) in Theorem 2.1 holds for all acts $f, g \in L(\succeq)$. The following statements are equivalent:*
(1) The preference relation $\succeq$ satisfies Symmetric Hedging.
(2) There exists a finitely additive probability measure $p_1 \in \mathcal{C}$ and a vector adjustment $p_{-1} \in \Delta^{\mathcal{V}}(\Omega,\Sigma)$ for p_1, taking values in a Banach space $(\mathcal{V}, \|\cdot\|_{\mathcal{V}})$, such that, for all acts $f, g \in L(\succeq)$,

$$f \succeq g \quad \Leftrightarrow \quad \int_\Omega u\circ f\; p_1(d\omega) - \Big\| \int_\Omega u\circ f\; p_{-1}(d\omega)\Big\|_{\mathcal{V}} \geq \int_\Omega u\circ g\; p_1(d\omega) - \Big\| \int_\Omega u\circ g\; p_{-1}(d\omega)\Big\|_{\mathcal{V}} \quad (7)$$

Furthermore, if (1) holds, then $\succeq$ satisfies Axiom (vi), Non-degeneracy, if and only if p_1 is unique and, for any vector adjustment p'_{-1} for p_1 that also satisfies Equation 7, there exists an isometric isomorphism $T : I(B;p_{-1},\Omega) \to I(B;p'_{-1},\Omega)$ such that $p'_{-1} = T\circ p_{-1}$.

Note that statement (2) may be reformulated as follows: there exists a probability measure p_1 and a vector-valued adjustment p_{-1} for p_1 such that, for all acts $f \in L(\succeq)$,

$$\min_{q\in\mathcal{C}} \int_\Omega u\circ f\; q(d\omega) = \int_\Omega u\circ f\; p_1(d\omega) - \Big\| \int_\Omega u\circ f\; p_{-1}(d\omega)\Big\|_{\mathcal{V}}. \quad (8)$$

The second part of Theorem 3.2 indicates the uniqueness properties of the VEU representation. In particular, the baseline probability measure p_1 is unique, and the vector measure p_{-1} that determines the adjustments associated with events (indicator functions) and acts are unique up to an isometric isomorphism—a norm-preserving, linear homeomorphism.

In the present setting, two measures p_{-1} and p'_{-1} related by such a transformation T convey the same decision-theoretically relevant information. First, they encode the same adjustments to baseline probabilities, because the transformation T is norm-preserving. Second, they encode the same information concerning the extent to which ambiguity about an event "cancels out" or "adds to" ambiguity about another, because T preserves the additivity properties of p_{-1}. More generally, they convey the same information on the effects of utility smoothing, because T also preserves the linearity properties of the integral operator $I(\cdot;p_{-1},\Omega)$, as well as norms on the entire subspace $I(B;p_{-1},\Omega)$. Example 2 below exemplifies the first two points; see also Example 7 in Section 4 for further details on utility smoothing.

Example 2 Refer to the VEU preferences described in the Introduction in the context of the Ellsberg Paradox. In particular, recall that $\Omega = \{\rho, \gamma, \beta\}$ and $p_{-1}(\rho) = 0$, $p_{-1}(\gamma) = -p_{-1}(\beta) = \frac{1}{3}$; the fact that the latter two adjustments have opposite signs indicates that ambiguity about γ and β "cancels out." Note that $I(B; p_{-1}, \Omega)$ is the entire real line.

Now consider a measure p'_{-1} taking values in $\mathbb{R}^2$, endowed with the Euclidean norm, denoted by $\|\cdot\|'$. If p'_{-1} is isometrically isomorphic to p_{-1}, then $\|p'_{-1}(\rho)\|' = 0$ and $\|p'_{-1}(\beta)\|' = \|p'_{-1}(\gamma)\|'_2 = \frac{1}{3}$. This implies that $p'_{-1}(\rho) = (0,0)$—the zero vector in $\mathbb{R}^2$; moreover, since additivity is preserved, $p'_{-1}(\gamma) = -p'_{-1}(\beta)$. Therefore, p'_{-1} encodes exactly the same information about γ and β as p_{-1}: the two events are equally ambiguous, but their ambiguities "cancel out".

Geometrically, $p'_{-1}(\gamma)$ and $p'_{-1}(\beta)$ are opposite points on a circle centered at the origin with radius equal to $\frac{1}{3}$, and $I(B; p'_{-1}, \Omega)$ is a line through the origin. Indeed, $\mathbb{R}^2$-valued measures isometrically isomorphic to p_{-1} are in one-to-one correspondence with such lines. This intuitively suggests that ambiguity in the Ellsberg Paradox is really "one-dimensional", regardless of the particular vector representation one chooses.

Incidentally, the example emphasizes that the measures p_{-1} and p'_{-1} in the statement of Theorem 3.2 may take values in different Banach spaces.

The proof of Theorem 3.2 also identifies a "canonical" vector adjustment. In the statement of the following Corollary, the notation $C(K)$ indicates the Banach space of bounded, continuous real functions defined on a compact set K, endowed with the supremum norm.

Corollary 3.3 *If (1) holds, the set function $p_{\mathcal{C}} : \Sigma \to \mathbb{R}^{\mathcal{C}}$ defined by $p_{\mathcal{C}}(E)(q) = p_1(E) - q(E)$ for all $E \in \Sigma$ and $q \in \mathcal{C}$ is a $C(\mathcal{C})$-valued adjustment for the probability measure p_1 whose existence is asserted in (2); moreover, p_1 and $p_{\mathcal{C}}$ jointly satisfy Equation 7 for all $f, g \in L(\succeq)$.*

4 Discussion and Extensions

The first part of this section discusses key aspects of Definition 2 and Symmetric Hedging, emphasizing their relationship with the Ambiguity Aversion axiom. The second collects additional results and examples.

4.1 Interpretation of Symmetric Hedging

As indicated in the previous section, the Schmeidler [18] - Gliboa and Schmeidler [11] Ambiguity Aversion axiom formally postulates a preference for "utility smoothing". A possible interpretation, related to a notion of ambiguity, runs as follows: if the decision maker cannot confidently assess the relative likelihood of certain events, she may wish to reduce perceived[4] differences in outcomes contingent upon their realization. The motivation for the Symmetric Hedging axiom builds upon these ideas.

The following example clarifies the suggested interpretation of Ambiguity Aversion, and will be referred to throughout this section.

Example 3 Consider the version of the Ellsberg paradox described in the Introduction. Recall that the decision maker is informed that a ball will be drawn from an urn containing 90 balls, of which 30 are red and the remainder green and blue, in unspecified proportions. The set of states is $\Omega = \{\rho, \gamma, \beta\}$, and the set of available prizes is $X = \{\$0, \$10\}$. Consider further the acts defined in Table 1.

Act	ρ	γ	β
g	\$0	\$10	\$0
b	\$0	\$0	\$10
$\frac{1}{2}g + \frac{1}{2}b$	\$0	$\frac{1}{2}\$10 + \frac{1}{2}\$0$	$\frac{1}{2}\$10 + \frac{1}{2}\$0$
$\bar{g}$	\$10	\$0	\$10
$\bar{b}$	\$10	\$10	\$0
$\frac{1}{2}\bar{g} + \frac{1}{2}\bar{b}$	\$10	$\frac{1}{2}\$10 + \frac{1}{2}\$0$	$\frac{1}{2}\$10 + \frac{1}{2}\$0$

Table 1: Acts in a three-color urn problem

As noted in the Introduction, the act g represents a bet "on" a green draw, whereas the act $\bar{g}$ is a bet "against" a green draw. The acts b and $\bar{b}$ are similarly interpreted.

The preferences $g \sim b$ and $\frac{1}{2}g + \frac{1}{2}b \succ g$ are consistent with Ambiguity Aversion; for instance, a MEU decision maker with beliefs represented by the class of priors $\mathcal{C} = \{p \in \Delta(\Omega) : p(\rho) = \frac{1}{3}\}$ and utility function u such that $u(\$10) > u(\$0)$ would display these preferences. Note that the latter also admit a VEU representation, with beliefs as indicated in the Introduction.

Note that any mixture $\alpha g + (1-\alpha)b$ reduces outcome variability across the elementary events γ and β, relative to the acts g or b. In particular, choosing $\alpha = \frac{1}{2}$ eliminates all variability. A preference for mixtures of g and b is then consistent with the idea that, since the ratio of green to blue balls is unspecified, the decision maker might favor acts whose ultimate outcome is less (or not at all) dependent upon the realization of γ rather than β.

Finally, note that the same observations hold for mixtures of the acts $\bar{g}$ and $\bar{b}$. Indeed, loosely speaking, a mixture of $\bar{g}$ and $\bar{b}$ reduces outcome variability relative to $\bar{g}$ or $\bar{b}$ to the same extent

[4] An act may yield objectively distinct, but subjectively equivalent lotteries contingent upon the realization of different events. In this case, there is obviously no scope for "utility smoothing".

as the same mixture of g and b reduces it relative to g or b. Consistent with this observation, a decision maker with beliefs as described above would express the preferences $\bar{g} \sim \bar{b}$ and $\frac{1}{2}\bar{g} + \frac{1}{2}\bar{b} \succ \bar{g}$.

It may be verified that the pairs of acts $g, \bar{g}$ and $b, \bar{b}$ are complementary, and that these preferences satisfy Symmetric Hedging.

The following two subsections analyze the key ingredients appearing in the statement of the Symmetric Hedging axiom. In the interest of brevity, two ordered pairs of acts (f, g) and $(\bar{f}, \bar{g})$ will be deemed complementary if f and $\bar{f}$ and, respectively, g and $\bar{g}$ are complementary acts.

4.1.1 Complementary Acts

Definition 2 encodes two (related) requirements related to the idea that complementary pairs of acts provide the same opportunities for reducing perceived outcome variability by mixing.

First, note that every act h determines a partition of the state space Ω into maximal events on which h assigns subjectively equivalent lotteries.[5] It is clear from Definition 2 that complementary events determine the same partition of Ω. Thus, loosely speaking, up to indifference among lotteries, complementary acts are defined in terms of the same (ambiguous and/or unambiguous) events. In particular, since mixtures of complementary pairs of acts (with the same weights) are themselves complementary, they also induce the same partition of Ω.

Second, lotteries assigned at any two distinct states by complementary acts may be intuitively interpreted as being reverse-ordered, but equally "far apart" or "close". Consider two complementary acts f and $\bar{f}$, and two distinct states ω, ω'; then the ranking

$$\frac{1}{2}f(\omega) + \frac{1}{2}\bar{f}(\omega) \sim \frac{1}{2}f(\omega') + \frac{1}{2}\bar{f}(\omega')$$

may be interpreted as imposing a restriction on the final outcome of two operations: starting with the mixture $\frac{1}{2}f(\omega) + \frac{1}{2}\bar{f}(\omega)$, (1) $f(\omega)$ is replaced with $f(\omega')$, and (2) $\bar{f}(\omega)$ is replaced with $\bar{f}(\omega')$. Since the decision maker is indifferent between the initial and final lotteries, either each replacement leaves the decision maker indifferent, or one leaves her worse off but the other exactly compensates for this. In this sense, $f(\omega)$ and $f(\omega')$ are as "far apart" as $\bar{f}(\omega)$ and $\bar{f}(\omega')$.

It is easy to see that, if the affine function u represents the decision maker's preferences among lotteries, then, for every $\omega \in \Omega$, $u(\bar{f}(\omega)) = k - u(f(\omega))$ for some constant $k \in \mathbb{R}$ (cf. the proof of Lemma A.2 in the Appendix). In other words, in utility terms, complementary acts are the negative of each other, possibily translated by a constant.[6]

[5] Formally, define a relation $=_h$ on Ω by letting $\omega =_h \omega'$ if and only if $h(\omega) \sim h(\omega')$. The partition alluded to in the text is the collection of equivalence classes of $=_h$.

[6] This property implies that complementary acts induce the same partition of Ω.

In particular, complementary acts and mixtures of complementary pairs of acts are therefore subject to the same perceived outcome variability, i.e. the same variability of utilities.

These observations formalize the heretofore informal assertion that complementary pairs provide the *same opportunities for utility smoothing*: their mixtures induce the same partition of Ω in the sense described above; furthermore, they are subject to the same variability of utilities.

In Example 3, it is easy to verify that (g, b) and $(\bar{g}, \bar{b})$ are complementary pairs, and that these acts and their mixtures exhibit the two features discussed here. For instance, note that the acts g and $\bar{g}$ both induce the partition $\{\{\gamma\}, \{\rho, \beta\}\}$, which intuitively, reflects the fact that the "source" of ambiguity in the evaluation of g and $\bar{g}$ is the same: namely, the fact that the ratio of green to blue balls is unspecified. Equal reduction of outcome variability is immediately apparent upon inspection of Table 1.

4.1.2 Symmetric Hedging

As mentioned in Section 3, the Symmetric Hedging axioms requires that pairs of acts that provide the same utility smoothing opportunities be perceived as providing an equally effective "hedge against ambiguity." This subsection comments on two aspects of this requirement.

First, the sense in which complementary pairs of acts provide the same utility smoothing opportunities has been made precise in Subsection 4.1.1. In particular, outcomes assigned by mixtures of complementary pairs at any two distinct states are *equidistant*, i.e. equally far apart in utility terms, albeit ordered in reverse fashion.

It seems natural to require that "hedging against ambiguity" be equally effective also for pairs of acts whose mixtures assign equidistant and *equally ordered* outcomes at any pair of states. Formally, for any two acts f, f', say that f' is a *traslation* of f if and only if, for every pair of states ω, ω',

$$\frac{1}{2}f(\omega) + \frac{1}{2}\bar{f}(\omega') \sim \frac{1}{2}\bar{f}(\omega) + \frac{1}{2}f(\omega').$$

The interpretation is as in Subsection 4.1.1: replacing (1) $f(\omega)$ with $\bar{f}(\omega)$ and (2) $\bar{f}(\omega')$ with $f(\omega')$ leaves the decision maker indifferent, so either (1) and (2) individually leave her indifferent, or one leaves her worse off and the other exactly compensates for this. Thus, in particular, if $f(\omega) \succ \bar{f}(\omega)$, then also $\bar{f}(\omega') \prec f(\omega')$.

One may then consider the following property:

TRASLATION-INVARIANT HEDGING. For all acts $f, g, f', g' \in L$, and for all $\gamma \in [0, 1]$: if f (resp. g) is a translation of f' (resp. g') and there exist a real number $\alpha \in [0, 1]$ and lotteries $y, y' \in Y$ such that

$$\alpha\Big(\gamma f + (1-\gamma)g\Big) + (1-\alpha)y \sim \alpha\Big(\gamma y(f) + (1-\gamma)y(g)\Big) + (1-\alpha)y' \tag{9}$$

then

$$\alpha\Big(\gamma f' + (1-\gamma)g'\Big) + (1-\alpha)y \sim \alpha\Big(\gamma y(f') + (1-\gamma)y(g'))\Big) + (1-\alpha)y' \tag{10}$$

The following remark states[7] that, in fact, MEU preferences automatically satisfy the above requirement. Hence, by imposing Symmetric Hedging, one ensures that *all* pairs of acts that provide equal opportunities for utility smoothing be perceived as providing an equally effective "hedge against ambiguity".

Remark 4.1 *Consider a preference relation $\succeq$ and assume that, for some utility function $u : Y \to \mathbb{R}$ and set of priors $\mathcal{C} \subset \Delta(\Omega, \Sigma)$, Equation 3 holds on $L(\succeq)$. Then:*

(a) for any two acts $f, f' \in L$, f' is a traslation of f if and only if there exists $k \in \mathbb{R}$ such that, for all $\omega \in \Omega$, $u(f'(\omega)) = u(f(\omega)) + k$;

(b) $\succeq$ satisfies Traslation-Invariant Hedging.

The second observation concerns the relationship between Symmetric Hedging and the notion of "ambiguity" implicit in the Ambiguity Aversion axiom.

The characterization of MEU preferences in [11] does not directly employ or refer to any formal definition of "ambiguous event". Rather, a preference for utility smoothing is interpreted as a consequence of (and reaction to) perceived ambiguity. Under this interpretation, violations of Ambiguity Neutrality may be taken to *reveal* the presence of ambiguity.[8]

Pursuing this line of inquiry to develop a definition of ambiguity within the Anscombe-Aumann framework is beyond the scope of this paper.[9] However, I suggest, by means of an example, that Symmetric Hedging may facilitate inferences about the presence of ambiguity based on violations of Ambiguity Neutrality. Equivalently, the following example suggests that the interpretation of the Ambiguity Aversion axiom may be less transparent when Symmetric Hedging is violated.

Example 4 [Another three-color urn problem] As in Example 3, let $\Omega = \{\rho, \gamma, \beta\}$ and consider the acts defined in Table 1. The composition of the urn and the beliefs of the decision maker are *not* the same as in Example 3: for now, I shall focus on preferences, and indicate their MEU representation later. Thus, assume:

$$g \sim b, \qquad \frac{1}{2}g + \frac{1}{2}b \sim g \tag{11}$$

[7] The proof is available upon request.

[8] As will be argued in Subsection 4.2.2, this seems legitimate in the current decision-theoretic setting, which incorporates the assumption that the decision maker behaves like an expected utility maximizer when evaluating lotteries.

[9] The interested reader is referred to Epstein [6], Epstein and Zhang [7] and Ghirardato and Marinacci [9] for behavioral definitions of ambiguity which do not rely on the existence of objective lotteries.

$$\bar{g} \sim \bar{b}, \qquad \frac{1}{2}\bar{g} + \frac{1}{2}\bar{b} \succ \bar{g} \tag{12}$$

Recall that the acts $g, \bar{g}$ and $b, \bar{b}$ are complementary, and hence provide the same opportunities for utility smoothing. Note that the preferences in Equation 11 are consistent with Ambiguity Neutrality, and hence with EU maximization. Thus, on the sole basis of these preferences, one is tempted to conclude that the events γ and β are *not* ambiguous. However, the preferences in Equation 12 clearly lead to the opposite conclusion, invoking the arguments given in Example 3.

These preferences have a MEU representation, whereby the set of priors is $\mathcal{C} = \Delta(\Omega)$. However, since they clearly violate Symmetric Hedging, they do not admit a VEU representation.

I emphasize that Example 4 does *not* indicate or imply that the preferences indicated above are "unintuitive"; it merely suggests that the "utility smoothing" interpretation of Ambiguity Aversion may be appropriate only for certain types of MEU preferences.

4.2 Additional Results

4.2.1 Symmetric Hedging and α-MEU Preferences

Apart from its role in the characterization of VEU preferences, the Symmetric Hedging axiom has implications for a class of representations which generalize the MEU model.

It has been observed that the "$\frac{1}{2}$-MEU functional" $f \mapsto \frac{1}{2}\min_{q\in\mathcal{C}} \int fq(d\omega) + \frac{1}{2}\max_{q\in\mathcal{C}} \int fq(d\omega)$ intuitively conveys a notion of "ambiguity neutrality". Ghirardato, Klibanoff and Marinacci [8] show that the above functional is indeed additive if the class $\mathcal{C}$ of priors is the convex hull of two probability measures, and provide an example to demonstrate that this is not the case in general.

This paper completes their analysis by providing, as a by-product, a complete characterization of the class of priors for which the $\frac{1}{2}$-MEU functional is additive.[10]

Specifically, if a preference relation $\succeq$ admits a MEU representation via a utility function u and a class of priors $\mathcal{C}$, then it satisfies Symmetric Hedging *if and only if* the "$\frac{1}{2}$-MEU" functional is additive (cf. Lemma A.2).

More generally, for $\alpha \in [\frac{1}{2}, 1]$, consider the "$\alpha$-MEU" functional $f \mapsto M_\alpha(f; \mathcal{C}, E)$ defined by

$$\forall f \in B, \quad M_\alpha(f; \mathcal{C}, E) = \alpha \min_{q\in\mathcal{C}} \int_E fq(d\omega) + (1-\alpha) \max_{q\in\mathcal{C}} \int_E fq(d\omega).$$

Again, α appears to provide a measure of "ambiguity aversion," for a *fixed* class of priors representing the decision maker's beliefs. This intuition can be made precise using the results in the present paper.

[10] This implies that the Remark following Example 2 in Section 4.2 of [8] is inaccurate.

Suppose that the MEU (i.e. "1-MEU") preferences generated by $\mathcal{C}$ satisfy Symmetric Hedging, so the $\frac{1}{2}$-MEU functional can be represented as the integral with respect to a probability measure $p \in \mathcal{C}$. Then straightforward calculations show that, for every $\alpha \in (\frac{1}{2}, 1]$ the class of priors

$$\mathcal{C}_\alpha = \{q' \in \Delta(\Omega, \Sigma) : q' = (2\alpha - 1)q + 2(1-\alpha)p \text{ for some } q \in \mathcal{C}\}$$

provides a 1-MEU representation of the α-MEU functional. That is, omitting the region of integration for notational simplicity, for every $f \in B$, $M_\alpha(f; \mathcal{C}) = M_1(f; \mathcal{C}_\alpha)$. Since the sets $\mathcal{C}_\alpha$ are ordered by inclusion ($\alpha > \alpha'$ implies $\mathcal{C}_\alpha \supset \mathcal{C}_{\alpha'}$), the index α provides a measure of the extent the decision maker deviates from EU maximization. Indeed, according to the comparative definitions of ambiguity and ambiguity aversion in Epstein [6] (provided the class $\mathcal{C}$ satisfies the appropriate restrictions on the set of unambiguous events) and in Ghirardato and Marinacci [9], the index α can be interpreted as an ambiguity aversion parameter—assuming, of course, that the underlying description of beliefs $\mathcal{C}$ is given. An analogous construction can be carried out for $\alpha < \frac{1}{2}$.

If the $\frac{1}{2}$-MEU functional is not additive, the above is not true in general, as one can easily show by considering a set Ω containing more than two elements and the class $\mathcal{C} = \Delta(\Omega, \Sigma)$.[11]

4.2.2 VEU and Probabilistic Sophistication

The VEU model allows for departures from probabilistic sophistication, as defined by Machina and Schmeidler [15, 16]; for instance, it accommodates the modal preferences in the Ellsberg example. This subsection provides a result and related comments on the possibility and consequences of probabilistic sophistication for VEU preferences. However, a preliminary observation is in order.

Recall that, by definition, the beliefs of a probabilistically sophisticated decision maker can be represented by a probability measure $\mu \in \Delta(\Omega, \Sigma)$, whereas her risk attitudes are captured by a preference functional W defined on the set of simple lotteries over final prizes; in the current setting, this is the set Y. Finally, note that every act f induces a lottery y_f via the measure μ[12]; then, probabilistically sophisticated preferences admit the following representation: $f \succeq g$ if and only if $W(y_f) \geq W(y_g)$.[13]

[11] The following converse also holds: if the α-MEU functionals are superlinear, and hence generate preferences satisfying Ambiguity Aversion, for $\alpha > \frac{1}{2}$, and are sublinear (hence generate preferences satisfying Ambiguity Appeal) for $\alpha < \frac{1}{2}$, then the $\frac{1}{2}$-MEU functional must be linear by a simple continuity argument. Therefore, the 1-MEU preferences satisfy Symmetric Hedging. Ghirardato, Maccheroni, Marinacci and Siniscalchi [10] analyze further properties of the classes of priors for which the $\frac{1}{2}$-MEU functional is linear.

[12] Formally, $y_f(x) = \int_\Omega f(\omega)(x) d\mu$ for every $x \in X$.

[13] Machina and Schmeidler [15] characterize probabilistically sophisticated preferences in a setting à la Savage, without objective lotteries; the authors also provide a simpler characterization in the Anscombe-Aumann setting in [16].

In particular, within the Anscombe-Aumann framework (cf. Machina and Schmeidler [16]), the functional W must represent preferences among acts *and* lotteries. But recall that, in this setting, the Gilboa-Schmeidler [11] axioms imply that a MEU decision maker behaves like an EU maximizer when choosing among lotteries. It follows that a MEU (hence, a VEU) decision maker is probabilistically sophisticated in the Anscombe-Aumann environment if and only if her overall preferences are consistent with EU maximization.

This observation has two important implications for the present paper. First, deviations from EU maximization may be ascribed to violations of probabilistic sophistication, and may thus be interpreted as reflecting perceived ambiguity.

Second, in order to analyze probabilistic sophistication in the VEU model without incurring in either degeneracy or contradiction, it seems appropriate to consider a more general decision-theoretic framework.

Thus, throughout this subsection, I consider an arbitrary set of outcomes $\mathcal{X}$, and define acts to be Σ-measurable maps from Ω to $\mathcal{X}$. In this framework, a preference relation $\succeq$ on acts admits a VEU representation if there exist a baseline probability p_1, a vector adjustment p_{-1} for p_1, and a function $u : \mathcal{X} \to \mathbb{R}$ such that the functional $f \mapsto V(u \circ f; p_1, p_{-1}, \Omega)$ represents $\succeq$.[14]

It is still the case that, in this more general framework, the beliefs of a VEU decision maker are represented by the difference between the baseline probability of an event and the norm of its vector adjustment. More precisely, define a weak order $\succeq_\ell$ on $\Sigma \times \Sigma$ by stipulating that, for all events $E, F \in \Sigma$,

$$E \succeq_\ell F \quad \Leftrightarrow \quad (x, E; x', \Omega \setminus E) \succeq (x, F; x', \Omega \setminus F) \tag{13}$$

for some pair of outcomes $x, x' \in Y$ such that $x \succeq x'$; $(x, E; x', \Omega \setminus E)$ is the simple act that yields the outcome x if E obtains, and x' otherwise. It is immediate to verify that, for MEU preferences, $\succeq_\ell$ is well-defined, i.e. the specific choice of "good" and "bad" outcomes is irrelevant. Moreover, for VEU preferences, $E \succeq_\ell F$ if and only if $p_1(E) - \|p_{-1}(E)\| \geq p_1(F) - \|p_{-1}(F)\|$, where p_1 and p_{-1} are the appropriate baseline probability and vector adjustment.

The question then arises whether the likelihood ordering $\succeq_\ell$ may reflect probabilistically sophisticated *beliefs*: that is, whether there exists a probability measure $\mu \in \Delta(\Omega, \Sigma)$ such that $E \succeq_\ell F$ if and only if $\mu(E) \geq \mu(F)$.

The following example shows that even non-degenerate VEU beliefs may be probabilistically sophisticated in the present context.

Example 5 Assume that $\Omega = \{\omega_1, \omega_2, \omega_3\}$ and let μ be the uniform measure on Ω. Now define

[14] Although this paper does not provide an axiomatization of VEU preferences in this setting, I am confident that one can be obtained using the techniques in Klibanoff et al. [14].

$p_1 = \mu$ and an l_3^∞-valued[15] vector adjustment $p_{-1} = (p_2, p_3, p_4)$ by letting

$$p_{-1} = \begin{bmatrix} 2/9 & 2/9 & 0 \\ 0 & -2/9 & 2/9 \\ -2/9 & 0 & -2/9 \end{bmatrix}$$

The ij-th entry in the above matrix corresponds to the adjustment term $p_{j+1}(\omega_i)$. It is easy to verify that, for any $E \subset \Omega$, $p_1(E) - \|p_{-1}(E)\|_\infty = \mu^2(E)$. Hence, for any pair $E, F \subset \Omega$, $E \succeq_\ell F$ if and only if $\mu(E) \geq \mu(F)$.

Observe that, in the above example, the measure μ representing $\succeq_\ell$ coincides with the baseline probability. This is always the case if μ is convex-ranged[16]:

Proposition 4.1 *Consider a VEU preference relation $\succeq$ and let $p_1 \in \Delta(\Omega, \Sigma)$ be the corresponding baseline probability. If the induced likelihood ordering $\succeq_\ell$ is represented by a convex-ranged probability measure $\mu \in \Delta(\Omega, \Sigma)$, then $\mu = p_1$.*

Machina and Schmeidler [15] provide a characterization of probabilistically sophisticated non-expected utility preferences (and hence beliefs) in the setting of this subsection. In particular, the axioms they propose imply that the probability measure representing the decision maker's beliefs is convex-ranged. Thus, Proposition 4.1 applies to all preference relations satisfying the Machina-Schmeidler axioms.

This result underscores the central role of baseline probabilities in the VEU representation. Consider a VEU decision maker who is not also an EU maximizer, but nevertheless has probabilistically sophisticated beliefs. Then, in the setting of this subsection, vector adjustments reflect aspects of her risk attitudes that are not captured by the baseline EU term. Thus, it seems intuitively plausible to expect that, if the decision maker's beliefs can be separated from her risk attitudes, then the vector adjustments (which reflect the latter) should not alter the likelihood ordering of events induced by the baseline probabilities. Proposition 4.1 confirms that this is the case if probabilistic beliefs are convex-ranged.

4.3 Additional Examples

This subsection illustrates further features of vector adjustments by means of two simple examples.

As discussed in Examples 2 and 3, there is only one source of ambiguity in the Ellsberg paradox; the first example of this subsection instead deals with "multidimensional" ambiguity. Also, it

[15] That is, p_{-1} takes values in $\mathbb{R}^3$, endowed with the maximum norm: $\|p_{-1}(E)\|_\infty = \max_{i=2,3,4} |p_i(E)|$.

[16] A probability measure μ is convex-ranged if, for every event $E \in \Sigma$ such that $p(E) > 0$, and for every $\alpha \in (0, 1)$, there exists $A \in \Sigma$ such that $A \subset E$ and $\mu(A) = \alpha\mu(E)$.

demonstrates how vector adjustments may be constructed on the basis of assumptions about the decision maker's perception of ambiguity.

The second example provides a graphical representation of the classes of priors corresponding to different VEU beliefs, again in the setting of a three-color urn problem; moreover, it illustrates how different attitudes towards utility smoothing may be captured by choosing different types of norms.

Example 6 Consider the following variant of the Ellsberg model. As in Example 3, the state space is denoted by $\Omega = \{\rho, \gamma, \beta\}$. As in the Ellsberg Paradox, the ratio of green vs. blue balls is unspecified; however, the decision maker is now told that, out of 90 balls, between 15 and 45 are red. Intuitively, this introduces an additional source of ambiguity.

In order to represent these beliefs consistently with the VEU model, consider a uniform baseline probability p_1: that is, $p_1(\rho) = p_1(\gamma) = p_1(\beta) = \frac{1}{3}$. The problem is then to construct an appropriate vector adjustment p_{-1} for p_1, taking values in some Banach space $\mathcal{V}$. In particular, the following equalities reflect the preceding assumptions:

$$\|p_{-1}(\rho)\|_{\mathcal{V}} = \frac{1}{6}, \qquad \|p_{-1}(\gamma)\|_{\mathcal{V}} = \|p_{-1}(\beta)\|_{\mathcal{V}} = \frac{1}{3}. \tag{14}$$

Equation 14 immediately implies that $\mathcal{V} \neq \mathbb{R}$—that is, ambiguity is not "one-dimensional". To see this, note that if $p_{-1}(\gamma)$ and $p_{-1}(\beta)$ are scalars, then $\|p_{-1}(\gamma \cup \beta)\|_{\mathcal{V}} = |p_{-1}(\gamma \cup \beta)|$ is either 0 or $\frac{2}{3}$, whereas $\|p_{-1}(\rho)\|_{\mathcal{V}} = \| - p_{-1}(\gamma \cup \beta)\|_{\mathcal{V}} = \frac{1}{3}$.

Thus, consider $\mathcal{V} = \mathbb{R}^2$, endowed with the l^1 norm. It is convenient to write adjustment vectors explicitly: $p_{-1}(E) = (p_2(E), p_3(E))$ for any event $E \subset \Omega$; then $\|p_{-1}(E)\|_{\mathcal{V}} = |p_2(E)| + |p_3(E)|$.

The following, heuristic argument identifies an appropriate vector adjustment. Essentially, the scalar measure p_2 will be used to encode ambiguity about the relative likelihood of ρ vs. $\gamma \cup \beta$, whereas p_3 will reflect ambiguity about the relative likelihood of γ vs. β.

First, to ensure that $\|p_{-1}(\rho)\|_{\mathcal{V}} = \frac{1}{6}$ while providing a separate representation of the two sources of ambiguity, let $p_2(\rho) = \frac{1}{6}$ and $p_3(\rho) = 0$. Additivity and the normalization condition $p_{-1}(\Omega) = (0, 0)$ then imply that $p_2(\gamma) + p_2(\beta) = -\frac{1}{6}$ and $p_3(\gamma) + p_3(\beta) = 0$.

Next, ambiguity about the relative likelihood of γ vs. β intuitively affects the decision maker's willingness to bet on either of them in a symmetric fashion. Thus, it seems plausible to choose p_3 so that $|p_3(\gamma)| = |p_3(\beta)|$; for definiteness, assume that $p_3(\gamma) > 0$. Equation 14 then implies that $|p_2(\gamma)| = |p_2(\beta)|$ as well. Hence, since $p_2(\gamma \cup \beta) = -\frac{1}{6}$ and $|p_2(\gamma)| + |p_3(\gamma)| = \frac{1}{3}$, a vector adjustment p_{-1} satisfying Equation 14 is now fully identified:

$$p_{-1} = \begin{bmatrix} \frac{1}{6} & 0 \\ -\frac{1}{12} & \frac{3}{12} \\ -\frac{1}{12} & -\frac{3}{12} \end{bmatrix}$$

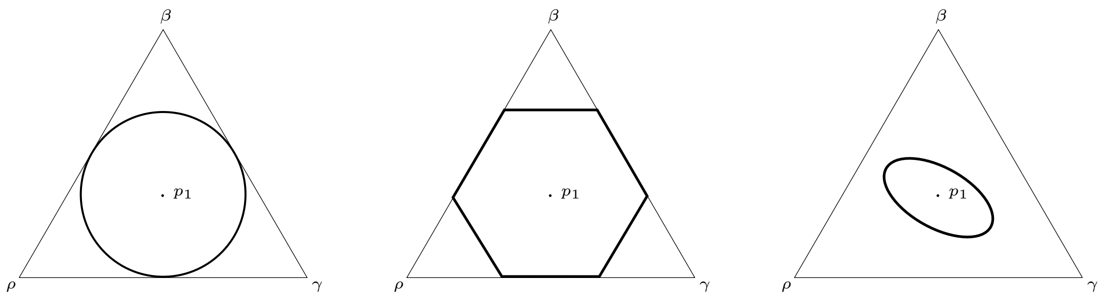

Case (a): l^2 norm

$$p_{-1} = \begin{bmatrix} \frac{1}{3\sqrt{2}} & \frac{1}{3\sqrt{2}} & 0 \\ -\frac{1}{3\sqrt{2}} & 0 & \frac{1}{3\sqrt{2}} \\ 0 & -\frac{1}{3\sqrt{2}} & -\frac{1}{3\sqrt{2}} \end{bmatrix}$$

Case (b): l^1 norm

$$p_{-1} = \begin{bmatrix} \frac{1}{6} & \frac{1}{6} & 0 \\ -\frac{1}{6} & 0 & \frac{1}{6} \\ 0 & -\frac{1}{6} & -\frac{1}{6} \end{bmatrix}$$

Case (c): l^2 norm

$$p_{-1} = \begin{bmatrix} \frac{1}{6} & 0 \\ -\frac{1}{6} & \frac{1}{6} \\ 0 & -\frac{1}{6} \end{bmatrix}$$

Figure 1: MEU Representations of VEU Beliefs ($p_1 = [\frac{1}{3}, \frac{1}{3}, \frac{1}{3}]'$)

where the ij-th element of the above matrix is $p_{j+1}(\omega_i)$ and $(\omega_1, \omega_2, \omega_3) = (\rho, \gamma, \beta)$.

Note that reversing the signs of one or both of the measures p_2 and p_3 yields a behaviorally equivalent (and isometrically isomorphic) vector adjustment for p_1. Indeed, any rotation of the vector p_{-1} about the origin yields an equivalent adjustment (cf. Example 2).

It should be observed that the family of vector adjustments with values in $(\mathbb{R}^2, l_2^1)$ is formally characterized by Equation 6 together with the linear restrictions imposed by additivity and normalization. However, the heuristic construction provided above is hopefully more revealing.

Example 7 Again, assume $\Omega = \{\rho, \gamma, \beta\}$; probability measures on Ω may be represented (using barycentric coordinates) as points in the 2-simplex.

Figure 1 plots the boundary of the set of priors corresponding to three different specifications of vector adjustments; in all three cases, the baseline probability is the uniform measure on Ω. As in the preceding example, the ij-th entry of the matrix p_{-1} is $p_{j+1}(\omega_i)$, and $(\omega_1, \omega_2, \omega_3) = (\rho, \gamma, \beta)$.

In Cases (a) and (b), the norm and vector measure are chosen to ensure that, for every event $E \subset \Omega$, the adjustment $\|p_{-1}\|_{\mathcal{V}}$ to $p_{-1}(E)$ is the same, i.e. $\frac{1}{3}$ for all sets consisting of one or two elements (recall the adjustment associated with the empty set and Ω is zero by definition).

However, it is interesting to point out that the different choice of norm has relevant behavioral implications. For instance, let the set of prizes be $X = \{\$0, \$10\}$ and let $u(\$0) = 0$, $u(\$10) = 1$. Consider the class of acts $f_x : \Omega \to Y$ defined by

$$f_x(\rho) = \$0, \quad f_x(\gamma) = x\$10 + (1-x)\$0, \quad f_x(\beta) = (1-x)\$10 + x\$0$$

for all $x \in [0,1]$. Thus, $u \circ f_x(\rho) = 0$, $u \circ f_x(\gamma) = x$ and $u \circ f_x(\beta) = 1 - x$.

Observe that $\alpha f_x + (1-\alpha) f_{x'} = f_{\alpha x + (1-\alpha)x'}$ for all $\alpha, x, x' \in [0,1]$; that is, the class of acts under consideration is closed under mixtures.

Simple calculations show that, in Case (a),

$$V(u \circ f_x) = \frac{1}{3} - \frac{1}{\sqrt{3}} \sqrt{x^2 - x + \frac{1}{3}}.$$

Since the above is a strictly concave function of x, for all numbers $\alpha, x, x' \in [0,1]$ one obtains the ordering $\alpha f_x + (1-\alpha) f_{x'} \succ \alpha y(f_x) + (1-\alpha) y(f_{x'})$. That is, the decision maker always displays a strict preference for reducing outcome variability (see also Remark 2.2).

On the other hand, in Case (b),

$$V(u \circ f_x) = \frac{1}{3} - \frac{1}{3} \max\Big(x, 1-x\Big).$$

Hence, for $\alpha \in [0,1]$ and either $x, x' \in [0, \frac{1}{2}]$ or $x, x' \in [\frac{1}{2}, 1]$, $\alpha f_x + (1-\alpha) f_{x'} \sim \alpha y(f_x) + (1-\alpha) y(f_{x'})$. The decision maker in Case (b) does not always display a strict preference for utility smoothing; it may be verified that she does so if the acts f_x and $f_{x'}$ are not comonotonic (Schmeidler [18]; see also Ghirardato, Klibanoff and Marinacci [8]), i.e. if $x < \frac{1}{2} < x'$.

The contrast between Cases (a) and (b) provides a further illustration of the notion of uniqueness of vector adjustments appearing in Theorem 3.2. The two vector measures indicated above generate the same adjustments for all baseline probabilities; however, adjustments associated with general acts may differ, and this violates the condition for behavioral equivalence stated in Theorem 3.2. More precisely, if p^a_{-1} and p^b_{-1} denote the vector adjustments in Cases (a) and (b) respectively, the subspaces $I(B; p^a_{-1})$ and $I(B; p^b_{-1})$ are *not* isometrically isomorphic.

The comparison also indicates that the VEU representation allows the modeler to characterize the decision maker's willingness to bet (via the map $E \mapsto p_1(E) - \|p_{-1}(E)\|$: see Subsection 4.2.2) separately from her attitudes towards utility smoothing (described by the choice of norm and range space $\mathcal{V}$).

References

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

[2] Aliprantis, C., and K. Border (1994): *Infinite Dimensional Analysis.* Springer Verlag, Berlin.

[3] Bartle, R.G. (1955): "A General Bilinear Vector Integral," *Studia Mathematica*, 15, 337-352.

[4] Dunford, N. and J. T. Schwartz (1957): *Linear Operators, Part I.* Wiley Interscience, New York.

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

[6] Epstein, L. (1999): "A Definition of Uncertainty Aversion," *Review of Economic Studies*, 66(3), 579-608.

[7] Epstein, L. and J. Zhang (2000): "Subjective Probabilities on Subjectively Unambiguous Events," *Econometrica*, forthcoming.

[8] Ghirardato, P., Klibanoff, P., and M. Marinacci (1998): "Additivity with Multiple Priors," *Journal of Mathematical Economics*, 30, 405-420.

[9] Ghirardato, P. and M. Marinacci (2000): "Ambiguity Made Precise: A Comparative Foundation," *Journal of Economic Theory*, forthcoming.

[10] Ghirardato, P., Maccheroni, F., Marinacci, M., and M. Siniscalchi (in preparation): "Symmetry and Ambiguity."

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

[12] Gilboa, I. and D. Schmeidler (1993): "Updating Ambiguous Beliefs," *Journal of Economic Theory*, 59, 33-49.

[13] Hansen, L.P., T. Sargent, and T. Tallarini (1999): "Robust Permanent Income and Pricing," *Review of Economic Studies*, 66(4), 873-908.

[14] Klibanoff, P., R. Casadesus-Masanell, and E. Ozdenoren (2000): "Maxmin-Expected Utility over Savage Acts with a Set of Priors," *Journal of Economic Theory*, 92(1), 35-65.

[15] Machina, M. and D. Schmeidler (1992): "A More Robust Definition of Subjective Probability," *Econometrica* 60, 4, 745-780.

[16] Machina, M. and D. Schmeidler (1995): "Bayes Without Bernoulli: Simple Conditions for Probabilistically Sophisticated Choice." *Journal of Economic Theory*, 67, 106-128.

[17] Savage, L.J. (1954), *The Foundations of Statistics*. Wiley, New York.

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

A Proofs

A.1 Additional Notation and Preliminaries

The following notation shall be employed to refer to the VEU and MEU functionals. Fix a function $f \in B(\Omega, \Sigma)$ and an event $E \in \Sigma$.

First, for any finitely additive probability measure $p_1 \in \Delta(\Omega, \Sigma)$ and vector adjustment $p_{-1} \in \Delta^{\mathcal{V}}(\Omega, \Sigma)$ for p_1, the vector-adjusted expectation of f over E with respect to (p_1, p_{-1}) is denoted by

$$V(f; p_1, p_{-1}, E) = \int_E f p_1(d\omega) - \left\| \int_E f p_{-1}(d\omega) \right\|_{\mathcal{V}}. \tag{15}$$

The first term in the right-hand side is a Dunford integral with respect to the finitely additive probability measure p_1. When referring to the functional V, the integrating measures and/or the region of integration will be omitted, if they can be inferred from the context. The main properties of the functional V are listed in Proposition A.1 below.

Next, consider a weak*-closed and convex set $\mathcal{C} \subset \Delta(\Omega, \Sigma)$ and a scalar $\alpha \in [0, 1]$. Define

$$M_\alpha(f; \mathcal{C}, E) = (1 - \alpha) \min_{q \in \mathcal{C}} \int_E f \, q(d\omega) + \alpha \max_{q \in \mathcal{C}} \int_E f \, q(d\omega). \tag{16}$$

The set of priors and region of integration will also be omitted, if they can be inferred from the context. The functional $M_1(\cdot; \mathcal{C}, E)$ will be denoted simply by $M(\cdot; \mathcal{C}, E)$; note that the latter is concave on B, and that

$$M_\alpha(f) = \alpha M(f) - (1 - \alpha) M(-f). \tag{17}$$

A.2 Properties of Vector-Adjusted Expectation and Proposition 3.1

Here and in the following, inequalities between functions are interpreted as pointwise inequalities.

Proposition A.1 *Fix a finitely additive probability measure $p_1 \in \Delta(\Omega, \Sigma)$ and a vector adjustment $p_{-1} \in \Delta^{\mathcal{V}}(\Omega, \Sigma)$ for p_1. The functional $V(\cdot; p_1, p_{-1}, \Omega) : B \to \mathbb{R}$ satisfies the following properties.*

(i) For any $c \in \mathbb{R}$, $V(c; p_1, p_{-1}, \Omega) = c$.

(ii) For any $f, g \in B$, $V(f + g; p_1, p_{-1}, \Omega) \geq V(f; p_1, p_{-1}, \Omega) + V(g; p_1, p_{-1}, \Omega)$.

(iii) For any $c \in \mathbb{R}$ and $f \in B$, $V(f + c; p_1, p_{-1}, \Omega) = V(f; p_1, p_{-1}, \Omega) + c$.

(iv) For any $c \in \mathbb{R}_+$ and $f \in B$, $V(cf; p_1, p_{-1}, \Omega) = cV(f; p_1, p_{-1}, \Omega)$.

(v) For any $f \in B$ such that $f \geq 0$, $V(f; p_1, p_{-1}, \Omega) \geq 0$; hence, $f \geq g$ implies $V(f; p_1, p_{-1}, \Omega) \geq V(g; p_1, p_{-1}, \Omega)$.

Observe that Condition (i) in Definition 1 is also necessary for Property (v) above to hold: if it fails for some $E \in \Sigma$, then $V(1_E; p_1, p_{-1}, \Omega) < 0$.

Proof: Part (i) is immediate. To simplify notation, let $E_J(f) = \int_\Omega f p_J(d\omega)$, for $J = 1$, $J = -1$; also write the norm on $\mathcal{V}$ simply as $\|\cdot\|$, and $V(\cdot)$ instead of $V(\cdot; p_1, p_{-1}, \Omega)$.

(ii) $V(f+g) = E_1(f+g) - \|E_{-1}(f+g)\| = E_1(f) + E_1(g) - \|E_{-1}(f) + E_{-1}(g)\| \geq E_1(f) + E_1(g) - \|E_{-1}(f)\| - \|E_{-1}(g)\| = V(f) + V(g)$, where the penultimate step follows from the triangle inequality.

(iii) $V(f+c) = E_1(f+c) - \|E_{-1}(f+c)\| = E_1(f) + c - \|E_{-1}(f) + E_{-1}(c)\| = E_1(f) - \|E_{-1}(f)\| + c = V(f) + c$, because, from (i), $E_{-1}(c) = 0$.

(iv) $V(cf) = E_1(cf) - \|E_{-1}(cf)\| = cE_1(f) - \|cE_{-1}(f)\| = cE_1(f) - c\|E_{-1}(f)\| = cV(f)$, where the penultimate step requires $c \geq 0$.

(v) First of all, note that the claim is true for all indicator functions: $V(1_E) = p_1(E) - \|p_{-1}(E)\| \geq 0$ by (i) in Definition 1. By (ii) and (iv) above, it is true for all simple functions. Hence, it holds for all uniform limits of simple functions, i.e. on all of B. ■

Proof of Proposition 3.1. By Proposition A.1, for any finitely additive probability measure p_1 and vector adjustment p_{-1}, the functional $V(\cdot; p_1, p_{-1}, \Omega)$ is monotonic, superadditive, homogeneous of degree 1, C-Linear (i.e. $V(f+c) = V(f) + V(c)$ for any function $f \in B$ and constant $c \in \mathbb{R}$), and satisfies $V(1) = 1$. Hence, by Lemma 3.5 in Gilboa and Schmeidler [11], there exists a closed and convex set $\mathcal{C}$ of finitely additive probability measures on Σ such that $V(f) = \min_{q \in \mathcal{C}} \int_\Omega f q(d\omega)$. ■

A.3 Remark 2.2

Proof: Suppose that $f \sim g$ and $\gamma f + (1-\gamma)g \succeq (\sim) \gamma y(f) + (1-\gamma)y(g)$. Then, by transitivity, $y(f) \sim y(g)$; by Independence on Y, $\gamma y(f) + (1-\gamma)y(g) \sim y(g) \sim g$; hence, by transitivity, Ambiguity Aversion (resp. Neutrality) holds.

Conversely, suppose $\succeq$ satisfies Ambiguity Aversion (resp. Neutrality). Then, by C-Independence, $f \sim y(f)$ and $g \sim y(g)$ imply $\gamma f + (1-\gamma)y(g) \sim \gamma y(f) + (1-\gamma)y(g)$ and $\gamma y(f) + (1-\gamma)g \sim \gamma y(f) + (1-\gamma)y(g)$, so $\gamma f + (1-\gamma)y(g) \sim \gamma y(f) + (1-\gamma)g$. Now, by Ambiguity Aversion (resp. Neutrality) we get

$$\frac{1}{2}\Big(\gamma f + (1-\gamma)y(g)\Big) + \frac{1}{2}\Big(\gamma y(f) + (1-\gamma)g\Big) \succeq (\sim) \gamma y(f) + (1-\gamma)g \sim \gamma y(f) + (1-\gamma)y(g),$$

which, upon rearrangement, becomes

$$\frac{1}{2}\Big(\gamma f + (1-\gamma)g\Big) + \frac{1}{2}\Big(\gamma y(f) + (1-\gamma)y(g)\Big) \succeq (\sim) \frac{1}{2}\Big(\gamma y(f) + (1-\gamma)y(g)\Big) + \frac{1}{2}\Big(\gamma y(f) + (1-\gamma)y(g)\Big),$$

and C-Independence yields the result. ■

A.4 Baseline Preferences and Symmetric Hedging

Throughout this subsection, suppose that $\succeq$ is a MEU preference relation, represented by the utility function u and the class of priors $\mathcal{C}$. Then a new preference relation $\succeq^b$ may be defined by letting

$$\forall f, g \in L(\succeq), \qquad f \succeq^b g \quad \Leftrightarrow \quad M_{\frac{1}{2}}(u \circ f; \mathcal{C}, \Omega) \geq M_{\frac{1}{2}}(u \circ g; \mathcal{C}, \Omega). \tag{18}$$

The superscript "b" indicates that $\succeq^b$ is the "baseline" preference that will yield the baseline probability measure in the VEU representation. Observe that $\succeq^b$ satisfies Axioms (i)-(iv), because the functional $M_{\frac{1}{2}}$ is monotonic, C-linear, positively homogeneous and continuous; moreover, $\succeq$ and $\succeq^b$ agree on L_c.

The following result contains the first key step in the proof that (1) implies (2) in Theorem 3.2. The second key step is contained in Lemma A.6.

Lemma A.2 *Assume that $\succeq$ is a MEU preference, represented by the class of priors $\mathcal{C}$ and the affine utility function u. Then it satisfies Symmetric Hedging iff $\succeq^b$ satisfies Ambiguity Neutrality.*

Proof: I begin with three preliminary observations.

Observation 1. Two acts $h, \bar{h}$ are complementary iff there exists $y_h \in Y$ such that, for every $\omega \in \Omega$, $u\left(\frac{1}{2}h(\omega) + \frac{1}{2}\bar{h}(\omega)\right) = \frac{1}{2}u(h(\omega)) + \frac{1}{2}u(\bar{h}(\omega)) = u(y_h)$, i.e. iff

$$u(\bar{h}(\omega)) = 2u(y_h) - u(h(\omega)). \tag{19}$$

Observation 2. For any two acts h, h', there exist $\alpha \in (0,1)$ and $y, y' \in Y$ such that

$$\alpha h + (1-\alpha)y \sim \alpha h' + (1-\alpha)y'$$

if and only if $M(u \circ h) - M(u \circ h') = \frac{1-\alpha}{\alpha}[u(y') - u(y)]$. To see this, observe that, by positive homogeneity and C-linearity of M, the above indifference relation holds if and only if

$$\alpha M(u \circ h) + (1-\alpha)u(y) = \alpha M(u \circ h') + (1-\alpha)u(y');$$

since $\alpha \in (0,1)$, the claim follows by rearranging terms.

Observation 3. For any pair of acts $h, h' \in L(\succeq)$, there exist $\alpha \in (0,1)$ and $y, y' \in Y$ such that $\alpha h + (1-\alpha)y \sim \alpha h' + (1-\alpha)y'$. The claim is trivially true if $h \sim h'$, so assume that $h \succ h'$ for definiteness and choose $y' \succ y$.[17] There exists $k \in \mathbb{R}_+$ such that $M(u \circ h) - M(u \circ h') = k[u(y') - u(y)]$; the claim now follows from Observation 2, by choosing $\alpha = \frac{1}{1+k} \in (0,1)$.

[17] This must be possible because $h \succ h'$.

Now consider four acts $f, g, \bar{f}, \bar{g}$ such that $f, \bar{f}$ and $g, \bar{g}$ are complementary, so Equation 19 holds for some lottery y_f and, respectively, y_g. By Observations 2 and 3, Symmetric Hedging requires that, for any such 4-tuple of acts, and for any $\gamma \in [0, 1]$,

$$M\Big(\gamma u \circ f + (1-\gamma) u \circ g\Big) - \Big(\gamma u(y(f)) + (1-\gamma) u(y(g))\Big) = \\ = M\Big(\gamma u \circ \bar{f} + (1-\gamma) u \circ \bar{g}\Big) - \Big(\gamma u(y(\bar{f})) + (1-\gamma) u(y(\bar{g}))\Big). \tag{20}$$

Recalling that $u(y(h)) = M(u \circ h)$ for any act h, this is equivalent to

$$M\Big(\gamma u \circ f + (1-\gamma) u \circ g\Big) - \Big(\gamma M(u \circ f) + (1-\gamma) M(u \circ g)\Big) = \\ = M\Big(\gamma u \circ \bar{f} + (1-\gamma) u \circ \bar{g}\Big) - \Big(\gamma M(u \circ \bar{f}) + (1-\gamma) M(u \circ \bar{g})\Big).$$

By Equation 19 and C-linearity of $M(\cdot)$, $M(u \circ \bar{h}) = M(2u(y_h) - u \circ h) = 2u(y_h) + M(-u \circ h)$ for any complementary pair of acts $h, \bar{h}$, with $y_h \sim \frac{1}{2}h + \frac{1}{2}h'$; moreover, $\gamma u \circ \bar{f} + (1-\gamma) u \circ \bar{g} = \gamma 2u(y_f) - \gamma u \circ f + (1-\gamma) 2u(y_g) - (1-\gamma) u \circ g$; thus, one obtains

$$M\Big(\gamma u \circ f + (1-\gamma) u \circ g\Big) - \Big(\gamma M(u \circ f) + (1-\gamma) M(u \circ g)\Big) = \\ = 2[\gamma u(y_f) + (1-\gamma) u(y_g)] + M\Big(-[\gamma u \circ f + (1-\gamma) u \circ g]\Big) + \\ -2[\gamma u(y_f) + (1-\gamma) u(y_g)] - \Big(\gamma M(-u \circ f) + (1-\gamma) M(-u \circ g)\Big).$$

Cancelling the common term $2[\gamma u(y_f) + (1-\gamma) u(y_g)]$, rearranging and dividing both sides by 2, one concludes that the 4-tuple of acts $f, g, \bar{f}, \bar{g}$ satisfies Equation 20 if and only if the pair f, g satisfies

$$M_{\frac{1}{2}}(\gamma u \circ f + (1-\gamma) u \circ g) = \gamma M_{\frac{1}{2}}(u \circ f) + (1-\gamma) M_{\frac{1}{2}}(u \circ g) \tag{21}$$

(cf. Equations 16 and 17).

The proof of the Lemma may now be completed. Suppose that $\succeq^b$ satisfies Ambiguity Neutrality; then Equation 21 holds for every pair of acts f, g, and hence Equation 20 holds for any 4-tuple of acts $f, g, \bar{f}, \bar{g}$, whenever $\bar{f}$ and $\bar{g}$ are complementary to f and g respectively. Hence, $\succeq$ satisfies Symmetric Hedging.

Conversely, suppose that $\succeq$ satisfies Symmetric Hedging. Note that, for any act $h \in L(\succeq)$, there exist $y', y'' \in Y$ such that $y' \succeq h(\omega) \succeq y''$ for every ω, so $u(u') \geq u(h(\omega)) \geq u(y'')$ for every ω. Define a new act $\bar{h}$ by letting

$$\forall \omega \in \Omega, \quad \bar{h}(\omega) = \frac{u(y') - u(h(\omega))}{u(y') - u(y'')} y' + \frac{u(h(\omega)) - u(y'')}{u(y') - u(y'')} y''; \tag{22}$$

observe that $\bar{h} \in L(\succeq)$ and, moreover, for $y = \frac{1}{2}y' + \frac{1}{2}y''$,

$$u(\bar{h}(\omega)) = u(y') + u(y'') - u(h(\omega)) = 2u(y) - u(h(\omega)),$$

i.e. h and $\bar{h}$ are complementary. Now consider any pair of acts $f, g \in L(\succeq) = L(\succeq^b)$. The acts $\bar{f}$ and $\bar{g}$ defined as in Equation 22 are complementary to f and g respectively. Since $\succeq$ satisfies Symmetric Hedging, the 4-tuple $f, g, \bar{f}, \bar{g}$ satisfies Equation 20, and hence the pair f, g satisfies Equation 21. Therefore $\succeq^b$ satisfies Ambiguity Neutrality. ■

Lemma A.2 implies that, if $\succeq$ is a MEU preference satisfying Symmetric Hedging, $\succeq^b$ is an EU preference, represented by some probability measure $p \in \Delta(\Omega, \Sigma)$. Lemma A.3 summarizes these conclusions, and adds that p is an element of the class of priors representing $\succeq$.

Lemma A.3 *Assume that the class of priors $\mathcal{C}$ and the affine utility function u provide a MEU representation of $\succeq$, and that $\succeq$ satisfies Symmetric Hedging. Then there exists a probability measure $p \in \Delta(\Omega, \Sigma)$ that, together with u, provides an EU representation of $\succeq^b$. Moreover, p is unique if and only if $\succeq$ satisfies Non-Degeneracy, in which case $p \in \mathcal{C}$.*

Proof: As noted above, $\succeq^b$ satisfies Axioms (i)-(iv) and agrees with $\succeq$ on L_c. By Lemma A.2, $\succeq^b$ satisfies Ambiguity Neutrality, so it admits an EU representation. In particular, there exists $p \in \Delta(\Omega, \Sigma)$ such that, for every $f \in L$,

$$\int_\Omega u \circ f p(d\omega) = \frac{1}{2} \min_{q \in \mathcal{C}} \int_\Omega u \circ f q(d\omega) + \frac{1}{2} \max_{q \in \mathcal{C}} \int_\Omega u \circ f q(d\omega). \tag{23}$$

Also, p is unique iff $\succeq^b$ satisfies Non-Degeneracy, hence iff $\succeq$ satisfies Non-Degeneracy. Moreover, $\int u \circ f p(d\omega) \geq M(u \circ f; \mathcal{C})$ for all acts $f \in L$ by Equation 23. To complete the proof of the Lemma, I claim that

$$\mathcal{M}(u, \mathcal{C}) \equiv \left\{ q \in \Delta(\Omega, \Sigma) \ : \ \forall f \in L, \int u \circ f q(d\omega) \geq M(u \circ f; \mathcal{C}) \right\} = \mathcal{C}.$$

To see this, note first that the set $\mathcal{M}(u, \mathcal{C})$ is weak*-closed and convex. Next, for every prior $q \in \mathcal{C}$, $\int u \circ f q(d\omega) \geq \min_{q' \in \mathcal{C}} \int u \circ f q'(d\omega) = M(u \circ f; \mathcal{C})$ for any act $f \in L$; hence, $\mathcal{C} \subset \mathcal{M}(u, \mathcal{C})$. This implies that, for any $f \in L$, $\min_{q \in \mathcal{M}(u,\mathcal{C})} \int u \circ f q(d\omega) \leq \min_{q \in \mathcal{C}} \int u \circ f q(d\omega) = M(u \circ f; \mathcal{C})$. But, by the definition of $\mathcal{M}(u, \mathcal{C})$, for any $f \in L$, $\min_{q \in \mathcal{M}(u,\mathcal{C})} \int u \circ f q(d\omega) \geq M(u \circ f; \mathcal{C})$; therefore, $\min_{q \in \mathcal{M}(u,\mathcal{C})} \int u \circ f q(d\omega) = M(u \circ f; \mathcal{C})$ for any $f \in L$.

Since $\mathcal{C}$ is the unique weak*-closed and convex set that provides a MEU representation of $\succeq$ because the latter satisfies Non-Degeneracy, $\mathcal{C} = \mathcal{M}(u, \mathcal{C})$, and the claim is proved. ■

To complete the analysis of baseline preferences, the next Lemma shows that, if $\succeq$ admits a VEU representation, then it satisfies Symmetric Hedging; this establishes that (2) implies (1) in Theorem 3.2. Moreover, the probability measure representing $\succeq^b$, whose existence is guaranteed by Lemma A.3, is actually the VEU baseline probability.

Lemma A.4 *Assume that the probability measure $p_1 \in \Delta(\Omega, \Sigma)$, the vector adjustment $p_{-1} \in \Delta^{\mathcal{V}}(\Omega, \Sigma)$ and the affine utility function u provide a VEU representation of $\succeq$, and the class of priors $\mathcal{C}$, together with u, provides the corresponding MEU representation. Then p_1 and u provide an EU representation of $\succeq^b$, and $\succeq$ satisfies Symmetric Hedging.*

Proof: Since $V(\cdot) = M(\cdot)$, for any function $b \in B$,

$$\begin{aligned} M_{\frac{1}{2}}(b) &= \frac{1}{2}M(b) - \frac{1}{2}M(-b) = \\ &= \frac{1}{2}\int bp_1(d\omega) - \frac{1}{2}\Big\| \int bp_{-1}(d\omega)\Big\| - \frac{1}{2}\int (-b)p_1(d\omega) + \frac{1}{2}\Big\| \int (-b)p_{-1}(d\omega)\Big\| = \\ &= \int bp_1(d\omega) \end{aligned}$$

using linearity of both integrals and the fact that $\| - v\| = \|v\|$ for $v \in \mathcal{V}$. Hence $\succeq^b$ admits an EU representation, and in particular satisfies Ambiguity Neutrality. By Lemma A.2, $\succeq$ satisfies Symmetric Hedging. ■

A.5 Norm representation of Adjustments

The results so far imply that a MEU preference relation $\succeq$ that satisfies Symmetric Hedging is represented by the difference between an EU term (which itself represents the baseline preference $\succeq^b$) and an adjustment term. Lemma A.6 below shows that the latter may be written as the norm of a vector integral, and hence completes the proof that (1) implies (2) in Theorem 3.2.

More precisely, for any weak*-closed and convex class $\mathcal{C}$ of priors, and event $E \in \Sigma$, define a functional $A(\cdot; \mathcal{C}, E) : B \to \mathbb{R}$ by

$$\forall f \in B, \quad A(f; \mathcal{C}, E) = \frac{1}{2}\left\{\max_{q\in\mathcal{C}} \int_E fq(d\omega) - \min_{q\in\mathcal{C}} \int_E fq(d\omega)\right\}. \tag{24}$$

Equation 23 implies that, if $\succeq$ is MEU preferences represented by the class $\mathcal{C}$ and the affine utility function u, then $M(u \circ f; \mathcal{C}) = \int_\Omega u \circ fp(d\omega) - A(u \circ f)$ for any act f. The next step is to extend this decomposition of the functional $M(\cdot)$ to all of B; the arguments are standard.

Lemma A.5 *In the setting of Lemma A.3, let $\mathcal{C}$ be the class of priors in the maxmin EU representation of $\succeq$. Then, for all $f \in B$ and $E \in \Sigma$,*

$$\begin{aligned} \int_E fp(d\omega) &= \frac{1}{2}\min_{q\in\mathcal{C}} \int_E fq(d\omega) + \frac{1}{2}\max_{q\in\mathcal{C}} \int_E fq(d\omega) \text{ and} \\ M(f; \mathcal{C}, E) &= \int_E fp(d\omega) - A(f; \mathcal{C}, E). \end{aligned}$$

Proof: Note first that if $c, d \in \mathbb{R}$ and $c \geq 0$, $A(cf + d) = cA(f)$.

By Equations (23) and (24), it is clear that the claim, with $E = \Omega$, holds for all *functions* $f' \in B$ such that, for some $f \in L$, $f' = u \circ f$. In particular, it holds on $u \circ L_0 = \{b \in B_0 : \exists f \in L_0, b = u \circ f\}$, the set of images of simple acts.

Since L_0 includes all simple acts, $u \circ L_0$ is the set of simple functions whose range is contained in the range of u. Therefore, any simple function $b \in B_0 \setminus u \circ L_0$ can be written as $b = cb' + d$, for real numbers $c, d \in \mathbb{R}$ with $c > 0$ and $b' \in u \circ L_0$. Since $M(\cdot)$ and the expectation operator are positively homogeneous and C-Independent, and $A(\cdot)$ is positively homogeneous, it follows that the claim, with $E = \Omega$, holds on all of B_0. Therefore, the claim also holds on B_0 for arbitrary $E \in \Sigma$.

Finally, the equality can be extended to all of $B(\Omega, \Sigma)$ because $M(f) = \min_{q \in \mathcal{C}} \int f q(d\omega)$, $\max_{q \in \mathcal{C}} \int f q(d\omega) = -M(-f)$, and $\int f q(d\omega)$ are all continuous functionals on B (endowed with the sup norm topology) and B is the uniform closure of B_0. ■

Note that uniqueness of p implies uniqueness of $A(\cdot)$.

It remains to be shown that $A(\cdot)$ can be written as the norm of a vector-valued expectation. Define the vector-valued map $p_{\mathcal{C}} : \Sigma \to \mathbb{R}^{\mathcal{C}}$ by

$$\forall E \in \Sigma, \qquad p_{\mathcal{C}}(E) = \{p(E) - q(E)\}_{q \in \mathcal{C}}. \tag{25}$$

Recall that $C(K)$ denotes the Banach space of bounded, continuous real functions on a compact space K, endowed with the sup norm; moreover, $\mathcal{C}$ is weak*-compact.

The next Lemma shows that $p_{\mathcal{C}}$ is a vector measure, yields the "right" adjustments for any $f \in B$, and satisfies Definition 1. As noted above, this completes the proof that (1) implies (2) in Theorem 3.2.

Lemma A.6 *Consider the setting of Lemma A.5. The set function $p_{\mathcal{C}}$ is a $C(\mathcal{C})$-valued measure on (Ω, Σ). Moreover, for all $f \in B$ and $E \in \Sigma$,*

$$A(f; \mathcal{C}, E) = \left\| \int_E f p_{\mathcal{C}}(d\omega) \right\|_{C(\mathcal{C})}.$$

Finally, the vector measure $p_{\mathcal{C}}$ satisfies $p_{\mathcal{C}}(\emptyset) = p_{\mathcal{C}}(\Omega) = 0^K$, and for every $E \in \Sigma$, $p(E) \geq \|p_{\mathcal{C}}(E)\|_{C(\mathcal{C})}$.

Proof: Recall that, in the topology of weak* convergence (cf. [4], §IV.13.17), $q^n \to q$ implies that $q^n(E) \to q(E)$ for all $E \in \Sigma$. Hence, if $q^n \to q$ in $\mathcal{C}$, $p_{\mathcal{C}}(E)(q^n) = p(E) - q^n(E) \to p(E) - q(E) = p_{\mathcal{C}}(E)(q)$ for all $E \in \Sigma$.[18] In other words, $p_{\mathcal{C}}(E) \in C(\mathcal{C})$ for all $E \in \Sigma$; moreover, $p_{\mathcal{C}}$ is finitely

[18] Also recall that the space of bounded additive measures on (Ω, Σ) is endowed with the metric induced by the total variation norm.

additive: hence, it is a $C(\mathcal{C})$-valued measure.

Now fix $f \in B$ and let $F = \{\int_E f p(d\omega) - \int_E f q(d\omega)\}_{q\in\mathcal{C}}$. Note that $F : \mathcal{C} \to \mathbb{R}$ is continuous and bounded, i.e. $F \in C(\mathcal{C})$. I claim that $\int_E f p_{\mathcal{C}}(d\omega) = F$.

The claim is true if $f \in B_0$: if $f = \sum_{k=1}^K x_k 1_{E_k}$, then

$$\begin{aligned}\int_E f p_{\mathcal{C}}(d\omega) &= \sum_{k=1}^K \{p(E_k \cap E) - q(E_k \cap E)\}_{q\in\mathcal{C}} \cdot x_k = \\ &= \left\{\sum_{k=1}^K p(E_k \cap E)x_k - \sum_{k=1}^K q(E_k \cap E)x_k\right\}_{q\in\mathcal{C}} = \\ &= F.\end{aligned}$$

Now consider a sequence $\{f^n\}_{n\geq 0}$ of simple functions converging monotonically and uniformly to f from below.[19] Then $\int_E f^n q(d\omega) \uparrow \int_E f q(d\omega)$ for every $q \in \mathcal{C}$, and $\int_E f^n p(d\omega) \uparrow \int_E f p(d\omega)$. Also, for every n, define $F^n = \{\int_E f^n p(d\omega) - \int_E f^n q(d\omega)\}_{q\in\mathcal{C}} = \int_E f^n p_{\mathcal{C}}(d\omega)$, and note that $F^n \in C(\mathcal{C})$. Moreover, $F^n \to F$ pointwise. It must be shown that convergence is actually uniform, so that $F^n \to F$ in the norm topology on $C(\mathcal{C})$.

To see this, note first that the integrals $\int_E f^n q(d\omega)$, for $n \geq 0$, as well as $\int_E f q(d\omega)$, are continuous functions of q, and convergence of the former to the latter is monotonic on the compact set $\mathcal{C}$; therefore, by Dini's Theorem (Aliprantis and Border [2], §2.63), $\int_E f^n q(d\omega) \to \int_E f q(d\omega)$ uniformly in q; thus, $\sup_{q\in\mathcal{C}} \left|\int_E f^n q(d\omega) - \int_E f q(d\omega)\right| \to 0$ as $n \to \infty$.

Next, note that

$$\begin{aligned}&\sup_{q\in\mathcal{C}} \left|\int_E f p(d\omega) - \int_E f q(d\omega) - \int_E f^n p(d\omega) + \int_E f^n q(d\omega)\right| \leq \\ &\leq \left|\int_E f p(d\omega) - \int_E f^n p(d\omega)\right| + \sup_{q\in\mathcal{C}} \left|\int_E f q(d\omega) - \int_E f^n q(d\omega)\right|.\end{aligned}$$

The first term in the r.h.s. vanishes as $n \to \infty$ by definition of the scalar integral; the second vanishes by the argument just given. Therefore, $F^n \to F$ in the norm topology on $C(\mathcal{C})$, and the claim is proved.

Now fix an event $E \in \Sigma$ and a function $f \in B$. Choose $q_{f,\min} \in \arg\min_{q\in\mathcal{C}} \int_E f q(d\omega)$ and $q_{f,\max} \in \arg\max_{q\in\mathcal{C}} \int_E f q(d\omega)$. Then $A(f;\mathcal{C},E) = \frac{1}{2}\{\int_E f q_{f,\max}(d\omega) - \int_E f q_{f,\min}(d\omega)\}$. Also, by

[19] E.g. let $f_{\inf} = \inf_{\omega\in\Omega} f(\omega)$ and $f_{\sup} = \sup_{\omega\in\Omega} f(\omega)$, both finite because $f \in B$. Assume that $f_{\sup} > f_{\inf}$ (otherwise take $f^n = f$, constant) and let $\hat{f}(\omega) = (f_{\sup} - f_{\inf})^{-1}(f(\omega) - f_{\inf})$ for every $\omega \in \Omega$. For every $n \geq 0$ and $k = 0, \ldots, 2^n - 1$, let $E_{k,n} = \{\omega : \hat{f}(\omega) \in [k2^{-n}, (k+1)2^{-n})\}$, and $E_{n,n} = \hat{f}^{-1}(1)$. Then let $\hat{f}^n = \sum_{k=0}^{2^n} 1_{E_{n,k}} k2^{-n}$. Note that, for every ω, if $\hat{f}^n(\omega) = k2^{-n}$, then $\hat{f}^{n+1}$ equals either $2k2^{-(n+1)} = k2^{-n}$ or $(2k+1)2^{-(n+1)} > k2^n$; moreover, $0 \leq \hat{f}(\omega) - \hat{f}^n(\omega) \leq 2^{-n}$. Hence, letting $f^n = f_{\inf} + (f_{\sup} - f_{\inf})\hat{f}^n$ yields the required sequence.

Lemma A.5, $\int_E fp(d\omega) = \frac{1}{2}\{\int_E fq_{f,\max}(d\omega) + \int_E fq_{f,\min}(d\omega)\}$; therefore,

$$\int_E fp(d\omega) - \int_E fq_{f,\min}(d\omega) = \int_E fq_{f,\max}(d\omega) - \int_E fp(d\omega) = A(f;\mathcal{C},E).$$

Now pick any $q \in \mathcal{C}$. If $\int_E fq(d\omega) \le \int_E fp(d\omega)$, then $\int_E fp(d\omega) - \int_E fq(d\omega) \le \int_E fp(d\omega) - \int_E fq_{f,\min}(d\omega)$; if instead $\int_E fq(d\omega) \ge \int_E fp(d\omega)$, then $\int_E fq(d\omega) - \int_E fp(d\omega) \le \int_E fq_{f,\max}(d\omega) - \int_E fq(d\omega)$.

Hence, $\sup_{q\in\mathcal{C}} \left|\int_E fp(d\omega) - \int_E fq(d\omega)\right| = \|\int_\Omega fp_{\mathcal{C}}(d\omega)\|_{C(\mathcal{C})}$ is attained by $q_{f,\min}$ and $q_{f,\max}$, and its value is $A(f;\mathcal{C},E)$, as required.

Finally, it is clear that $p_{\mathcal{C}}(\emptyset) = p_{\mathcal{C}}(\Omega) = 0^{\mathcal{C}}$. Also, $p(E) - \|p_{\mathcal{C}}(E)\|_{C(\mathcal{C})} = \int_\Omega 1_E dp - A(1_E;\mathcal{C},\Omega) = M(1_E;\mathcal{C},\Omega) \ge 0$. ■

The last Lemma in this section establishes the uniqueness claim in Theorem 3.2.

Lemma A.7 *Fix two vector-valued measures μ and μ' on (Ω,Σ), with values in $(\mathcal{V},\|\cdot\|)$ and $(\mathcal{V}',\|\cdot\|')$ respectively. If, for all $f \in B$, $\|\int_\Omega f\mu(d\omega)\| = \|\int_\Omega f\mu'(d\omega)\|'$, then there exists an isometric isomorphism $T : I(B;\mu,\Omega) \to I(B;\mu',\Omega)$ such that $\mu' = T\circ\mu$.*

Proof: Denote by I and I' the range spaces in the statement. Both are normed linear spaces. Define $T : I \to I'$ by letting, for every $v \in I$, $T(v) = \int f\mu'(d\omega)$ if $f \in B$ is such that $\int f\mu(d\omega) = v$.

To see that T is well-defined, note that if $f, g \in B$ satisfy $\int f\mu(d\omega) = \int g\mu(d\omega) = v$, then $\|\int(f-g)\mu(d\omega)\| = 0 = \|\int(f-g)\mu'(d\omega)\|'$; hence, $\int g\mu'(d\omega) = \int f\mu'(d\omega)$.

To see that T is one-one, suppose $v = \int f\mu(d\omega) \neq \int g\mu(d\omega) = w$; then $0 \neq \|\int(f-g)\mu(d\omega)\| = \|\int(f-g)\mu'(d\omega)\|'$, so $T(v) = \int f\mu'(d\omega) \neq \int g\mu'(d\omega) = T(w)$.

T is linear: if $\alpha,\beta \in \mathbb{R}$ and $v = \int f\mu(d\omega)$, $w = \int g\mu(d\omega)$, then $\alpha v + \beta w = \int(\alpha f + \beta g)\mu(d\omega)$, so by definition $T(\alpha v + \beta w) = \int(\alpha f + \beta g)\mu'(d\omega) = \alpha T(v) + \beta T(w)$.

By assumption, for every $v \in I$, $\|T(v)\|' = \|v\|$; therefore, T is a bounded, hence continuous linear operator. Therefore, T is an isometric isomorphism between I and I'. Furthermore, for every $E \in \Sigma$, $T(p_{-1}(E)) = \int 1_E p'_{-1}(d\omega) = p'_{-1}(E)$. ■

A.6 Proposition 4.1.

Since $\succeq_\ell$ is represented by μ and also by the map defined by $E \mapsto p_1(E) - \|p_{-1}(E)\|$, there exists an increasing function $g : [0,1] \to [0,1]$ such that $p_1(E) - \|p_{-1}(E)\| = g(\mu(E))$ for all events E. Hence, recalling that $p_{-1}(\Omega\setminus E) = -p_{-1}(E)$,

$$\frac{1}{2}g(\mu(E)) + \frac{1}{2}[1 - g(1-\mu(E))] = p_1(E) \tag{26}$$

for all events $E \in \Sigma$. Now, since μ is convex-ranged, for any integer n there exists a partition $\{E_1^n, \ldots, E_n^n\}$ of Ω such that $\mu(E_j^n) = \frac{1}{n}$ for all $j = 1, \ldots, n$. By (26), we have

$$\frac{1}{2} g\left(\frac{1}{n}\right) + \frac{1}{2}\left[1 - g\left(\frac{n-1}{n}\right)\right] = p_1(E_j^n), \quad j = 1, \ldots, n.$$

Adding up these equalities and dividing by n yields

$$\frac{1}{2} g\left(\frac{1}{n}\right) + \frac{1}{2}\left[1 - g\left(\frac{n-1}{n}\right)\right] = \frac{1}{n}.$$

Comparing the last two equalities yields $p_1(E_j^n) = \frac{1}{n} = \mu(E_j^n)$ for all $j = 1, \ldots, n$. This implies that, for every event E such that $\mu(E)$ is rational, $p_1(E) = \mu(E)$.

To extend this equality to arbitrary events, note that, for every event E such that $\mu(E) > 0$ and number $r < \mu(E)$, since μ is convex-ranged, there exists $L \subset E$ such that $\mu(L) = \frac{r}{\mu(E)}\mu(E) = r$. Similarly, for every event E such that $\mu(E) < 1$ and number $r > \mu(E)$, there exists an event $U \supset E$ such that $\mu(U) = r$: to see this, note that $\mu(\Omega \setminus E) > 0$ and $1 - r < \mu(\Omega \setminus E)$, so there exists $L \subset \Omega \setminus E$ such that $\mu(L) = 1 - r$; hence, $U = \Omega \setminus L$ has the required properties.

Now consider sequences of rational numbers $\{\ell_n\}_{n \geq 0} \subset [0, 1]$ and $\{u_n\}_{n \geq 0} \subset [0, 1]$ such that $\ell_n \uparrow \mu(E)$ and $u_n \downarrow \mu(E)$; by the preceding argument, for every $n \geq 1$ there exist sets $L_n \subset E \subset U_n$ such that $\mu(L_n) = \ell_n$ and $\mu(U_n) = u_n$. It was shown above that $p_1(L_n) = \mu(L_n)$ and $p_1(U_n) = \mu(U_n)$; moreover, $L_n \subset E \subset U_n$ implies that $p_1(L_n) \leq p_1(E) \leq p_1(U_n)$. Therefore, $p_1(E) = \mu(E)$, as required. ■