Journal of Risk and Uncertainty, 12:189–207 (1997)

Dynamically Consistent Preferences with Quadratic Beliefs

JÜRGEN EICHBERGER
The University of Saarland

SIMON GRANT
The Australian National University

Abstract

This article characterizes a family of preference relations over uncertain prospects that (a) are dynamically consistent in the Machina sense and, moreover, for which the updated preferences are also members of this family and (b) can simultaneously accommodate Ellsberg- and Allais-type paradoxes.

Replacing the "mixture independence" axiom by "mixture symmetry," proposed by Chew, Epstein, and Segal (1991) for decision making under objective risk, and requiring that for some partition of the state space the agent perceives ambiguity and so prefers a randomization over outcomes across that partition (*proper uncertainty aversion*), preferences can be represented by a (proper) quadratic functional. This representation may be further refined to allow a separation between the quantification of beliefs and risk preferences that is closed under dynamically consistent updating.

Key words: dynamic consistency, quadratic utility, quadratic beliefs

JEL Classification: D81

1. Introduction

Despite a growing empirical, experimental, and theoretical challenge over the past two decades, subjective expected utility theory (SEUT) remains the dominant model for analyzing decision making under uncertainty. This is particularly the case for dynamic choice situations where uncertainty is only gradually resolved through time and individuals may be afforded at different points in time the opportunity to revise their plans of action. Various researchers have demonstrated that, by imposing certain consistency requirements on an individual's dynamic choice, the individual's preferences must conform to SEUT (see, e.g., Hammond, 1988; Karni and Schmeidler, 1991). These have provided a "normative" defense for the contention that SEUT is the only "rational" model for choice under uncertainty (see, e.g., Howard, (1992).

Machina's (1989) response has been to point out that such consistency requirements essentially embody a *consequentialist* approach that does not allow uncertainty that has already been borne to influence preferences for uncertainty that is still to be resolved. This entails a separability of preferences across mutually exclusive states that is entirely inappropriate, in Machina's view, for preferences that do not conform to SEUT. For non-SEUT preferences, dynamic consistency can be satisfied tautologically if choice at any stage of a dynamic choice problem is determined by the conditional (or updated) preference relation appropriately defined from the original unconditional preference relation. McClennen (1990) refers to such dynamically consistent choice as "resolute."[1]

By itself, dynamic consistency in the Machina sense has little bite for "rational" decision making under uncertainty, except for the usual ordering requirement of the unconditional preference relation. Generally, however, what is deemed by the analyst to be the unconditional universe of the decision problem facing the individual may not correspond to a fundamental initial period for the decision maker. That is to say, the unconditional preferences may themselves be viewed as conditional preferences that have resulted from the resolution of some earlier (but unmodeled) uncertainty. If, therefore, an individual's dynamically consistent behavior can be described by his or her "unconditional" preferences, then it seems natural to us to invoke Epstein and le Breton's (1993) principle, that the set of properties that characterize the individual's "unconditional" preferences should be inherited by all the conditional preference relations that can be derived from this "parent" relation. Hence, we shall require that if the "unconditional" preferences can be represented by a particular family of functionals, then all the derived conditional (updated) preference relations can also be represented by the same family of functionals[2].

We employ in this article the horse-race lottery framework of Anscombe and Aumann (1963) (hereafter, AA). A horse-race lottery is a compound lottery, because the (intermediate) outcome that an individual receives contingent upon the resolution of the subjective uncertainty is itself a lottery or gamble. The process that determines the ultimate outcome from the lottery prize is assumed to be "objective" in that the probability of any particular final outcome arising from the lottery prize is determined *independently* of the preferences of the individual. For a given set of outcomes, $\mathscr{X}$, and a given (finite) set of states $\mathscr{I}$, the set of horse-race lotteries is a mixture space similar in structure to the space of lotteries that features in the theory of decision making under risk. The restriction of the preference relation to the set of constant acts may be naturally identified with the "risk preferences" of the individual, because the set of constant acts is isomorphic to the set of probability distributions over the set of possible (final) outcomes.

AA characterize preferences that can be represented by a SEUT functional. In this case, beliefs over the likelihood of the subjectively uncertain states can be separated from risk preferences, the beliefs are probabilistically sophisticated (as they can be represented by a unique probability measure), and the risk preferences conform to expected utility theory. The well-known Ellsberg paradox, however, provides an instance where an individual's choices reveal a preference for "known" or "objective" risk over subjectively uncertain events and, as a consequence, his or her beliefs cannot be represented by a probability measure. Allais-type paradoxes on the other hand, involve choices where the probabilities are known, but the choice patterns observed demonstrate that risk preferences are not *linear in the probabilities* and, hence, do not conform to expected utility theory. In addition to dynamic and sequential consistency, we shall require that the unconditional preference relation of an individual can accommodate both types of paradoxes.

The purpose of this article is thus to characterize a family of preference relations over uncertain prospects that satisfy the following desiderata: Preferences (a) are dynamically consistent in the Machina sense and, moreover, the updated preferences are from the same class as the parent preferences; and (b) can simultaneously accommodate Ellsberg-type paradoxes; that is, preferences are *ambiguity averse*, and Allais-type paradoxes, that is, risk preferences are not linear in the probabilities.

We primarily accomplish this by replacing the "mixture independence" axiom employed by AA by its weaker cousin "mixture symmetry" proposed by Chew, Epstein, and Sega (1991) in the context of decision making under objective risk in order to accommodate Allais-style paradoxes. Mixture symmetry requires that for any two indifferent horse-race lotteries, f and g, a probability mixture $\alpha f + (1 - \alpha)g$, with α in (0,1/2) is indifferent to a mixture $\beta f + (1 - \beta)g$, for some β in (1/2,1), though not necessarily indifferent to f and g. In addition, in order to accommodate Ellsberg paradox behavior, we require that for some partition of the state space the agent perceives ambiguity and so prefers objective mixing of final prizes across that partition (what we refer to as *proper uncertainty aversion*). With these axioms, we show that preferences can be represented by a (proper) quadratic functional. Furthermore, any updated (or conditional) preference relation derived from such a preference relation can also be represented by a quadratic functional[3].

We also show that this representation may be further refined to allow a separation between the quantification of beliefs and risk preferences that is closed under dynamically consistent updating. The quadratic functional obtained involves two probability measures defined on the set of events. Proper uncertainty aversion corresponds to these two probability measures differing.

We conclude the article with a reexamination of resolute choice. We note that consistent dynamic choice entails that individuals treat *experienced* ambiguity in exactly the same way as *anticipated* ambiguity. We employ an example based on a simple extension of Ellsberg's original set of decision problems to suggest that this might be an implausible way to model the behavior of ambiguity-averse individuals. Whether an individual's beliefs are quadratic or not, we contend that in modeling her choice behavior one should explicitly account for the structure of the resolution of uncertainty through time. In particular, one can specify the individual's conditional preferences at each decision node rather than defining a preference relation over the global set of acts.[4]

2. "Quadratic" preferences for horse-race lotteries

We adopt the Anscombe–Aumann setup for horse-race lotteries. Table 1 presents the basic concepts and introduces the notation.

Notice that $\mathcal{H}$ is a mixture space. For all $\lambda \in [0,1]$, and all $f, g \in \mathcal{H}$, $\lambda f + (1 - \lambda)g$ means the horse-race lottery that gives $\lambda f_s + (1 - \lambda)g_s$ for all $s \in \mathcal{S}$. We denote by $\succ$ the strict preference relation and by $\sim$ the indifference relation. Abusing notation we will write P, to refer to the horse-race lottery that gives the lottery prize P irrespective of the state that obtains. Similarly, we will write δ^x or x to refer to the horse-race lottery that

Table 1. Anscombe–Aumann "horse-race lottery" framework

$\mathcal{S} = \{s_1, \ldots, s_s\}$	a finite set of states.
$\mathcal{E} = \{A, B, \ldots\} = 2^S$	the set of events (the set of all subsets of $\mathcal{S}$).
$\mathcal{X} = \{x, y, z, \ldots\}$	a finite set of (final) outcomes or consequences.
$\Delta(\mathcal{X}) = \{P, Q, \ldots\}$	the set of probability distributions (hereafter, *lotteries*) on $\mathcal{X}$.
$\mathcal{H} = \{f, g, h, \ldots\}$	the set of mappings from $\mathcal{S}$ to $\Delta(\mathcal{X})$ (i.e. *horse-race lotteries*).
$\succsim$	a binary relation over ordered pairs of horse-lotteries representing the agent's preferences.

results in the final outcome x regardless of which state obtains. Sometimes, the following notation to describe a horse-race lottery will be convenient:

> For $E \in \mathcal{E}$, $f_E h$ is the horse-race lottery that gives for each state s the lottery prize f_s if s is in E, or the lottery prize h_s if s is not in E.

Using this notation we can define the set of *null events*, $\mathcal{N} \subset \mathcal{E}$, as follows: $E \in \mathcal{N}$ if and only if, for all $f, g, h \in \mathcal{H}$, $f_E h \sim g_E h$.

The following three axioms guarantee the existence of a nonzero function that represents the horse-race lotteries.

Order: $\succsim$ is a complete and transitive preference order.
Continuity: for all f, g, h in $\mathcal{H}$ such that $f \succ g \succ h$, the sets $\{\alpha \text{ in } [0, 1] \mid g \succ \alpha f + (1 - \alpha)h\}$ and $\{\alpha \text{ in } [0, 1] \mid \alpha f + (1 - \alpha)h \succ g\}$ are open.
Nondegeneracy: there exist acts f, g in $\mathcal{H}$ such that $f \succ g$.

Because the space of horse-race lotteries $\mathcal{H}$ can be viewed as a convex subset of a vector space, existence of a numerical representation follows from standard arguments. To get a further characterization, it is customary to assume that preferences satisfy some form of monotonicity and independence.

The basic monotonicity that we shall require is that preferences respect the partial ordering over lotteries of first-order stochastic dominance in all non-null states. That is, the ordinal ranking of a final outcome is invariant to the state in which it arises. Moreover, moving probability weight in a state's lottery prize from a less (unconditionally) preferred final outcome to a more (unconditionally) preferred final outcome can only improve the act in preference terms.

(Final) Prize Monotonicity:

for all outcomes x, y in $\mathcal{X}$, all lotteries P in $\Delta(\mathcal{X})$, all acts h in $\mathcal{H}$, all λ in $(0, 1]$ and all non-null events $E \subseteq \mathcal{S}$,

$$x \succsim y \Leftrightarrow [\lambda\delta^x + (1 - \lambda)P]_E h \succsim [\lambda\delta^y + (1 - \lambda)P]_E h.$$

One implication of Prize Monotonicity is that if there are two outcomes x and x', for which $x \sim x'$, then any occurrence of x' in any horse-race lottery h can be replaced by x without changing the desirability of that horse-race lottery. As prize monotonicity is assumed throughout this article, we will remove such "redundancy" in the preferences over horse-race lotteries by assuming that there does not exist a pair of (final) outcomes x and x' for which $x \sim x'$.

Prize Monotonicity is much weaker than the following state independence axiom that is usually imposed on horse-race lotteries.

Lottery State Independence:

for all lotteries P, Q in $\Delta(\mathscr{X})$, all acts h in $\mathscr{H}$, and all

$$\text{non-null events } E \subseteq \mathscr{I},$$
$$P \succsim Q \Leftrightarrow P_E h \succsim Q_E h.$$

This stronger axiom assumes that the individual is always willing to replace a risky lottery in an event by its certainty equivalent. It is too strong for our purposes as it entails that the certainty equivalent for a lottery prize depends neither on the event in which it obtains nor on what would have been received if that event had not obtained.

The other separability axiom that forms the "lynch-pin" of subjective expected utility in the Anscombe–Aumann horse-lottery framework is the mixture independence axiom.

Mixture Independence:

For every f, g and h in $\mathscr{H}$, and α *in* $[0,1]$,

$$f \succsim g \text{ implies } \alpha f + (1-\alpha)h \succsim \alpha g + (1-\alpha)h.$$

Stimulated by the Allais paradox, in the context of risk, one suggested direction for generalizing the expected utility framework has been to weaken mixture independence to hold only within an indifference set of lotteries. The corresponding axiom in the Anscombe–Aumann horse-race lottery framework is thus:

Betweenness:

For every f, g *in* $\mathscr{H}$ and α in $[0,1]$,

$$f \sim g \text{ implies } \alpha f + (1-\alpha)g \sim f.$$

Betweenness imposes "linearity" on the indifference sets, but unlike Mixture Independence does not require them to be "parallel." This particular weakening of Independence is not in accord with an intuitive interpretation of what it entails for an individual to distinguish between ambiguity and risk. Consider two horse-race lotteries f and g between which our agent is indifferent. Assume that both acts only assign final outcomes (i.e., degenerate lottery prizes) for each state. Let x_s (respectively, y_s) denote the outcome assigned by f (respectively, g) if state s is realized. Unless f and g always assign equally desirable outcomes for the realization of each state, then there will be at least one state s, for which $x_s \succ y_s$ and another state t, for which $y_t \succ x_t$. If we consider a mixture $\alpha f + (1 - \alpha)g$, for some α in (0,1/2), Prize Monotonicity implies that

$$x_s \succ \alpha\delta_s^x + (1 - \alpha)\delta_s^y \succ y_s \text{ and } y_t \succ \alpha\delta_t^x + (1 - \alpha)\delta_t^y \succ x_t.$$

Hence, mixing the horse-race lotteries f and g provides an insurance (or, as Schmeidler (1989) dubs it, *hedging*) element by averaging of the lottery prizes. Depending on a person's attitude toward such "insurance," one can distinguish between ambiguity and risk according to whether a person prefers the mixture to the original lotteries or vice versa.

Whichever it is, however, one might assume that for any probability mixture that places more weight on f, there is another probability mixture that places more weight on g such that the individual is indifferent between the two mixtures. Applying this property to all pairs of indifferent horse-race lotteries provides us with the horse-race lottery analogue of Chew, Epstein, and Segal's (1991) *mixture symmetry axiom*.

Mixture Symmetry:

For every pair of acts f,g in $\succsim$H, $f \sim g$ implies

for all $\alpha \in (0,1/2)$ there exists $\beta \in (1/2,1)$ such that

$$\alpha f + (1 - \alpha)g \sim \beta f + (1 - \beta)g.$$

One might reason further, that for each α the required β should be $(1 - \alpha)$; that is, "symmetric" mixtures of indifferent horse-race lotteries should themselves be indifferent. And indeed, as Chew, Epstein, and Segal (1991) have shown, given appropriate continuity and monotonicity assumptions for the preferences (Prize Monotonicity in our case) Mixture Symmetry is equivalent to the following axiom.

Strong Mixture Symmetry:

For every pair of acts f, g in $\mathcal{H}$, and $\alpha \in [0,1]$,

$$f \sim g \text{ implies } \alpha f + (1 - \alpha)g \sim (1 - \alpha)f + \alpha g.$$

As mentioned above, taking a mixture of two acts between which the agent is indifferent provides insurance across states. Clearly for a subjective expected utility maximizer, such averaging of utilities across states provides no benefit as the weighted average of the utility of the lottery prizes remains unchanged. For a person who perceives ambiguity in his or her assessed likelihood of the states, such a mixing averages the lottery prizes and thus may be valued (respectively, disliked) if the person dislikes (respectively, likes) the ambiguity. Of course, certain partitions of the state space may hold no ambiguity for the agent; thus for acts that are defined on such partitions, the individual views mixing in the same way as she or he views mixing purely risky objects.

One class of functionals that exhibit the mixture symmetry property is the class of quadratic functions.

$$V(f) = \sum_{s\in\mathcal{S}} \sum_{t\in\mathcal{S}} \sum_{x\in\mathcal{X}} \sum_{y\in\mathcal{X}} \psi(x,y,s,t)\, f_s(x) f_t(y).$$

However, any functional representing preferences that satisfies Betweenness also exhibits the mixture symmetry property and, as Chew, Epstein, and Segal (1991) point out, these two classes of functionals completely exhaust the functionals that satisfy mixture symmetry. Moreover, the only intersection of these two classes is the class of functionals that satisfy mixture independence.

In order to exclude the Betweenness class of functionals, we shall require that the agent perceives ambiguity for some partition of the state space and that preferences over acts defined on that partition exhibit aversion to such ambiguity.

For any two outcomes $x,y \in \mathcal{X}$ with $x > y$ and any $p \in [0,1]$, denote by $L(p)\!: = p \cdot \delta^x + (1 - p) \cdot \delta^y$ the lottery that yields the better outcome x with probability p and the worse outcome y with probability $(1 - p)$.

Proper uncertainty aversion:

For $A \in \mathcal{E}$, $x,y \in \mathcal{X}$ with $x > y$ and $p,q,r,s \in [0, 1]$ with $p > q$ and

$$r > s \text{ such that } L(p)_A L(q) \sim L(s)_A L(r), \text{ let } \lambda := \frac{(r - s)}{(p - q) + (r - s)},$$

$$\text{then } L(\lambda \cdot p + (1 - \lambda)s)_A L(\lambda \cdot q + (1 - \lambda)r) = L\left(\frac{pr - qs}{(p - q) + (r - s)}\right) \succcurlyeq L(p)_A L(q).$$

Moreover, there exists an event B $\in \mathcal{E}$ for which the strict preference holds.

The implication of Proper Uncertainty Aversion for $\succsim$ can be illustrated in figure 1. Consider two horse-race lotteries that are constant on a partition of the state space into subsets A and ¬A. Suppose that prizes are two-outcome lotteries on outcomes x and y. In figure 1, the distance along the horizontal (respectively, vertical) axis represents the

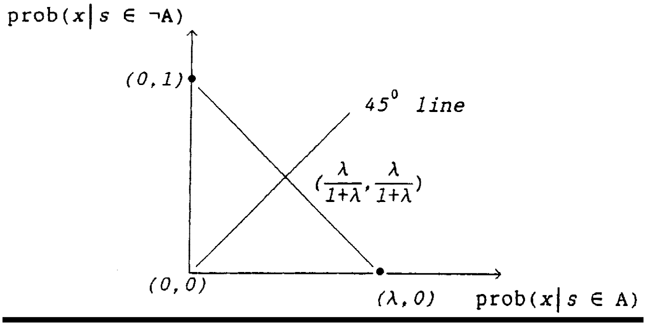

Figure 1.

probability that the lottery prize for event A (respectively, event ¬A) assigns to the final outcome x with the remaining probability weight assigned to outcome y. Hence, a horse-race lottery of the form

$$L(p)_A L(q) = [p\delta^x + (1 - p)\delta^y]_A[q\delta^x + (1 - q)\delta^y]$$

can be represented by a point (p,q) in figure 1.

The axiom states that if the individual is indifferent between $(\lambda,0)$ and $(0,1)$ then she weakly prefers (strictly prefers for some event B) the $[1/(1 + \lambda), \lambda/(1 + \lambda)]$ probability mixture of these two points to $(0,1)$. This mixture corresponds to the point $(\lambda/(1 + \lambda), \lambda/(1 + \lambda))$ in figure 1 for which there is no ambiguity. Proper ambiguity aversion entails that the indifference curve connecting $(0,1)$ and $(\lambda,0)$ cannot dip above the intersection of the 45° line with the chord that joins those two points. Moreover, it must dip below for some event B.

In conjunction with the other axioms (in particular, mixture symmetry), proper ambiguity aversion requires that the preference relation is convex, when restricted to acts that only differ on two disjoint events and that only assign probability weight to the same two outcomes on those two events. That is, the function that represents $\succsim$ restricted to lotteries such as the ones represented in figure 1, is quasi-concave (and for the particular case of event B is proper quasi-concave).

The following theorem follows as an almost immediate corollary of Chew, Epstein, and Segal's Theorem 5 (1991, p. 149). A formal proof is contained in the Appendix.

Representation Theorem. Let $\succsim$ satisfy Order, Continuity, Nondegeneracy, Ordinal State Independence, and Proper Uncertainty Aversion. Then $\succsim$ satisfies Mixture Symmetry if and only if it can be represented numerically by a proper quadratic utility function V of the form

$$V(f) = \sum_{s\in\mathcal{S}} \sum_{t\in\mathcal{S}} \sum_{x\in\mathcal{X}} \sum_{y\in\mathcal{X}} \psi\,(x,y,s,t)\, f_s(x) f_t(y) \tag{1}$$

with $\psi(x, y, s, t)$ increasing in x and y.

3. Dynamic consistency, separating beliefs from risk preferences and updating

In the introduction we suggested two desiderata for a representation of preferences. First, because it is derived in the Anscombe–Aumann framework, the representation should allow us to separate risk-preferences from beliefs. Note that risk-preferences are naturally given by the representation over constant acts. Second, dynamically consistent updates of preferences should leave the class of representations invariant. We will discuss in turn the implications of these two requirements on the quadratic representation (1) which we derived in section 2.

3.1. Dynamic consistency

Dynamic consistency requires that the representation in (1) be maintained after any event E that becomes known and for any act h previously chosen. In order to see the implication of such an assumption, define

$$\psi_1(x,s \mid E,h) := [\sum_{t \notin E} \sum_{y \in \mathcal{X}} \psi(x,y,s,t) h_t(y)] / |E|,$$

$$\psi_2(y,t \mid E,h) := [\sum_{s \notin E} \sum_{x \in \mathcal{X}} \psi(x,y,s,t) h_s(x)] / |E| \text{ and}$$

$$\psi_3(E,h) := [\sum_{s \notin E} \sum_{t \notin E} \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{X}} \psi(x,y,s,t) h_s(x) h_t(y)] / |E|^2,$$

where $|E|$ denotes the cardinality of the set E. For an observed event $E \subseteq S$ and a previously chosen act h, let

$$\hat{\psi}(x,y,s,t \mid E,h) := \psi(x,y,s,t) + \psi_1(x,s \mid E,h) + \psi_2(y,t \mid E,h) + \psi_3(E,h).$$

The following lemma shows the restrictions that follow from dynamically consistent updating.

Lemma 3.1. Suppose act h has been chosen and event E has been observed, then

$$V(f_E \mid E, h) = \sum_{s \in E} \sum_{t \in E} \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{X}} \hat{\psi}(x, y, s, t \mid E, h) f_s(x) f_t(y)$$

is the dynamically consistent updated representation of $V(f)$.

Proof. Noting that, for all acts f, $\sum_{s\in E}\sum_{x\in\mathcal{X}} f_s(x) = \sum_{t\in E}\sum_{y\in\mathcal{X}} f_t(y) = |E|$, one calculates easily:

$$
\begin{aligned}
V(f_E \mid E,h) &= \sum_{s\in E}\sum_{t\in E}\sum_{x\in\mathcal{X}}\sum_{y\in\mathcal{X}} \hat{\psi}(x,y,s,t \mid E,h) f_s(x) f_t(y) \\
&= \sum_{s\in E}\sum_{t\in E}\sum_{x\in\mathcal{X}}\sum_{y\in\mathcal{X}} [\psi(x,y,s,t) + \psi_1(x,s \mid E,h) \\
&\qquad + \psi_2(y,t \mid E,h) + \psi_3(E,h)] f_s(x) f_t(y) \\
&= \sum_{s\in E}\sum_{t\in E}\sum_{x\in\mathcal{X}}\sum_{y\in\mathcal{X}} \psi(x,y,s,t) f_s(x) f_t(y) + \\
&\quad [\sum_{t\in E}\sum_{y\in\mathcal{X}} f_t(y)] \cdot [\sum_{s\in E}\sum_{x\in\mathcal{X}} \psi_1(x,s \mid E,h) f_s(x)] + \\
&\quad [\sum_{s\in E}\sum_{x\in\mathcal{X}} f_s(x)] \cdot [\sum_{t\in E}\sum_{y\in\mathcal{X}} \psi_2(y,t \mid E,h) f_t(y)] + \\
&\quad [\sum_{s\in E}\sum_{t\in E}\sum_{x\in\mathcal{X}}\sum_{y\in\mathcal{X}} f_s(x) f_t(y)] \cdot \psi_3(E, h) \\
&= \sum_{s\in E}\sum_{t\in E}\sum_{x\in\mathcal{X}}\sum_{y\in\mathcal{X}} \psi(x,y,s,t) f_s(x) f_t(y) + \\
&\quad \sum_{s\in E}\sum_{t\in E}\sum_{x\in\mathcal{X}}\sum_{y\in\mathcal{X}} \psi(x,y,s,t) f_s(x)\, h_t(y) + \\
&\quad \sum_{s\in E}\sum_{t\in E}\sum_{x\in\mathcal{X}}\sum_{y\in\mathcal{X}} \psi(x,y,s,t)\, h_s(x) f_t(y) + \\
&\quad \sum_{s\in E}\sum_{t\in E}\sum_{x\in\mathcal{X}}\sum_{y\in\mathcal{X}} \psi(x,y,s,t)\, h_s(x)\, h_t(y) \\
&=: V(f_E h)
\end{aligned}
$$

■

Lemma 3.1 shows that the dynamically consistent update of the quadratic representation (1) will be a generalized quadradic form containing linear parts, ψ_1 and ψ_2, as well as a constant term, ψ_3. Any property of the representation ψ that is supposed to carry over to the updated representation $\hat{\psi}$ must therefore also hold for the linear terms ψ_1 and ψ_2.

3.2. Separation of beliefs and risk preferences

An important feature of subjective expected utility and Choquet expected utility in the Anscombe–Aumann framework is the separation between the quantification of beliefs and risk preferences in these representations. Because the constant acts induce a representation over lotteries, the evaluation of these lotteries can be separated from the weighting of these evaluations. Similarly, it follows from Chew, Epstein, and Segals (1991) that risk preferences can be determined in a natural way from the quadratic representation over constant acts.

One would therefore like to refine the representation in (1) to allow for a separation of risk preferences and beliefs. Moreover, such a separation should be retained for any updated preferences. As this subsection will demonstrate, the restriction to classes of representations that are closed under dynamically consistent updating imposes constraints on the type of separability.

Lemma 3.1 shows that any dynamically consistent separation of risk preferences from beliefs must respect the form of $\hat{\psi}(x,y,s,t \mid E,h)$. In particular, the separation of ψ must be such that beliefs and preferences over outcomes remain also separated in ψ_1 and ψ_2.

A natural separation of risk preferences and beliefs would separate the function ψ into a quadradic part over outcome pairs and a quadratic part over pairs of states: $\psi(x,y,s,t)\colon = \alpha(x,y)\cdot\beta(s,t)$. It is not difficult to check, however, that such a separation cannot separate beliefs from risk references in the linear terms ψ_1 or ψ_2, for example,

$$\begin{aligned}\psi_1(x,s \mid E,h)\colon &= [\sum_{t\notin E}\sum_{y\in\mathscr{X}} \alpha(x,y)\cdot\beta(s,t)\, h_t(y)] / \mid E \mid \\ &= [\sum_{t\notin E}\beta(s,t)\cdot\sum_{y\in\mathscr{X}}\alpha(x,y)\, h_t(y)] / \mid E \mid \\ &= [\sum_{t\notin E}\beta(s,t)\cdot K(x,t)]/\mid E \mid : m = L(x,s),\end{aligned}$$

with

$$K(x,t)\colon = \sum_{y\in\mathscr{X}}\alpha(x,y)\, h_t(y).$$

From the last line of this calculation, it follows that either $\alpha(x,y)$ or $\beta(s,t)$ must also be separable.

If $\alpha(x,y)\colon = u(x)\cdot v(y)$ holds, then one obtains

$$\psi_1(x,s \mid E,h) = u(x)\cdot[\sum_{t\notin E}\beta(s,t)\cdot E_{ht}v] / \mid E \mid \; : = u(x)\cdot M(s)$$

with $E_{ht}v := \sum_{y\in\mathscr{X}} v(y)\, h_t(y)$. Hence, for $\psi(x,y,s,t)\colon = u(x)\cdot v(y)\cdot\beta(s,t)$, the updated representation has the following form

$$\hat{\psi}(x,y,s,t \mid E,h) = u(x)\cdot v(y)\cdot\beta(s,t) + u(x)\cdot M(s) + v(y)\cdot N(t) + \gamma$$

where $N(t)\colon = [\sum_{s\notin E}\beta(s,t)\cdot E_{hs}u]/\mid E \mid$ and γ a constant. This separation will, however, not leave the functional form unchanged.

With this in mind, there is essentially one way how beliefs can be separated from risk preferences in a manner that is inherited by updated preferences:

$$\psi(x,y,s,t) := \phi(x,y) \cdot \mu(s) \cdot \nu(t).$$

W.l.o.g., one can assume that the belief functions μ and ν are normalized to satisfy $\sum_{s=1}^{s} \mu(s) = \sum_{t=1}^{s} \nu(t) = 1$ in this case. For any event E and any act h where $\mu(E) := \sum_{s\in E} \mu(s)$ and $\nu(E) := \sum_{t\in E} \nu(t)$ are both strictly greater than zero, let

$$K(x \mid E,h) := \sum_{t\notin E} \sum_{y\in\mathcal{X}} \phi(x,y)\nu(t)\, h_t(y),$$

$$L(y \mid E,h) := \sum_{s\notin E} \sum_{x\in\mathcal{X}} \phi(x,y)\mu(s)\, h_s(x),$$

$$H(E,h) := \sum_{s\notin E} \sum_{t\notin E} \sum_{x\in\mathcal{X}} \sum_{y\in\mathcal{X}} \phi(x,y)\mu(s)\, \nu(t)\, h_s(x)\, h_t(y).$$

One easily checks that

$$\psi_1(x,s \mid E,h) = [K(x \mid E,h) / \mid E \mid](\cdot)\mu(s),$$

$$\psi_2(y,t \mid E,h) = [L(y \mid E,h) / \mid E \mid](\cdot)\nu(t), \text{ and}$$

$$\psi_3(E,h) = H(E,h)/\mid E \mid^2.$$

This demonstrates the separability of ψ_1 and ψ_2 in beliefs about states and risk preferences for this representation. Renormalizing the beliefs,

$$\hat{\mu}_E(s) := \mu(s)/\mu(E), \textit{ and } \hat{\nu}_E(t) := \nu(t)/\nu(E),$$

and defining

$$\hat{\phi}(x,y \mid E,h) := \mu(E)\nu(E)\phi(x,y) + \mu(E)K(x \mid E,h) + \nu(E)L(y \mid E,h) + H(E,h),$$

the following lemma is immediate:

Lemma 3.3.

$$V(f_E \mid E,h) = \sum_{s\in E} \sum_{t\in E} \sum_{x\in\mathcal{X}} \sum_{y\in\mathcal{X}} [\hat{\phi}(x,y \mid E,h)\cdot\hat{\mu}_E(s)\cdot\hat{\nu}_E(t)]\, f_s(x)\, f_t(y).$$

Proof.

$$
\begin{aligned}
V(f_E \mid E,h) := {} & V(f_E h) \\
= {} & \sum_{s\in E}\sum_{t\in E}\sum_{x\in\mathcal{X}}\sum_{y\in\mathcal{X}} \phi(x,y)\mu(s)\nu(t)\, f_s(x) f_t(y) + \\
& \sum_{s\in E}\mu(s)\sum_{x\in\mathcal{X}} f_s(x)\,[\sum_{t\notin E}\sum_{y\in\mathcal{X}} \phi(x,y)\nu(t)\, h_t(y)] + \\
& \sum_{t\in E}\nu(t)\sum_{y\in\mathcal{X}} f_t(y)[\sum_{s\notin E}\sum_{x\in\mathcal{X}} \phi(x,y)\mu(s)\, h_s(x)] + \\
& \sum_{s\notin E}\sum_{t\notin E}\sum_{x\in\mathcal{X}}\sum_{y\in\mathcal{X}} \phi(x,y)\mu(s)\nu(t)\, h_s(x) h_t(y) \\
= {} & \sum_{s\in E}\sum_{t\in E}\sum_{x\in\mathcal{X}}\sum_{y\in\mathcal{X}} \phi(x,y)\mu(s)\nu(t)\, f_s(x) f_t(y) + \\
= {} & \sum_{s\in E}\mu(s)\sum_{x\in\mathcal{X}} f_s(x)\, K(x \mid E,h) + \sum_{t\in E}\nu(t)\sum_{y\in\mathcal{X}} f_t(y)\, L(y \mid E,h) + H(E,h) \\
= {} & \sum_{s\in E}\sum_{t\in E}\sum_{x\in\mathcal{X}}\sum_{y\in\mathcal{X}} [\phi(x,y)\mu(E)\nu(E) + K(x \mid E,h)\mu(E) + L(y \mid E,h)\nu(E) + \\
& H(E,h)]\,[\hat{\mu}_E(s)\hat{\nu}_E(t)]\, f_s(x)\, f_t(y) \\
= {} & \sum_{s\in E}\sum_{t\in E}\sum_{x\in\mathcal{X}}\sum_{y\in\mathcal{X}} [\hat{\phi}(x,y \mid E,h)\hat{\mu}_E(s)\hat{\nu}_E(t)]\, f_s(x) f_t(y).
\end{aligned}
$$

■

For this representation, beliefs and their updates are uniquely defined because μ, ν, $\hat{\mu}$ and $\hat{\nu}$ are probability measures. The valuations of output ϕ and $\hat{\phi}$ are unique up to affine transformations. Note that this representation characterizes beliefs by two probability distributions over states. As the following section will show, two probability distributions are sufficient to represent beliefs that are ambiguity averse.

4. Accommodation of Allais- and Ellsberg-style paradoxes

The separable form $\psi(x,y,s,t) = \phi(x,y)\mu(s)\nu(t)$ that was shown in the section above to lead to updated risk preferences from the same class, is also flexible enough to accommodate Allais- and Ellsberg-style paradoxes.

4.1. Common ratio paradox

Consider the following simple variant of the Allais (Common Ratio) paradox:

- lottery I pays \$70 for sure

and

- lottery II pays \$100 with probability 0.8 and nothing with probability 0.2.

Most subjects in experimental studies when faced with such a choice express a preference for the "risk-free" lottery I. If they were expected-utility maximizers, one would then conclude that $[u(70) - u(0)]/[u(100) - u(0)]$ was greater than 0.8. Many of the same subjects when presented with the choice between

- lottery III that pays \$70 with probability 0.25 and nothing with probability 0.75

and

- lottery IV that pays \$100 with probability 0.2 and nothing with probability 0.8,

tended to prefer the lottery with the higher expected payoff, that is lottery IV over lottery III. But this would suggest that if they were expected-utility maximizers, $[u(70) - u(0)]/[u(100) - u(0)]$ was less than 0.8. Such a preference pattern is clearly inconsistent with expected utility theory, but can be readily accommodated by a quadratic functional of the form above.

For example, let

$$\phi(x,y) = \frac{1}{2}[u(x)v(y) + v(x)u(y)]$$

where $u(100) = 100$, $u(70) = 95$, $u(0) = 10$ and $v(y) = y$. Then we have (rounded to the nearest integer)

$$V(I) = 6650 > V(II) = 6560 \textit{ and } V(III) = 547 < V(IV) = 560.$$

4.2. Ellsberg paradox

With the following simple variant of the original Ellsberg (1961) paradox, one can argue against the claim in the Savage subjective expected utility model of decision making under uncertainty that beliefs of decision makers can be represented by probability distributions. Consider an urn that contains 20 red balls and 40 balls that are either black or white. Faced with the choice between

- lottery A paying \$100 on a draw of a red ball and nothing otherwise

and

- lottery B paying \$100 on a draw of a white ball and nothing otherwise,

most subjects preferred lottery A. One could conclude from this choice that these decision makers assume that the probability of drawing a white ball is smaller than the probability of drawing a red ball, which equals 1/3. The same subjects, however, given the choice between

- lottery C paying \$100 for a draw of either a white or a black ball and nothing otherwise

and

- lottery D paying \$100 for a draw of a red or a black ball and nothing otherwise,

tended to prefer lottery C over lottery D, which would suggest that they assume that drawing a white ball is more likely than drawing a red ball, which occurs with probability 1/3. Such decisions are clearly inconsistent with probabilistic beliefs. Ellsberg explained this change in the assessed likelihood of states by a preference for unambiguous beliefs, because the probability of drawing a red ball is known to be 1/3 and the probability of a white or black ball is objectively 2/3. All other events do not have objective probabilities assigned to them.

Quadratic beliefs as they are proposed in this article can accommodate the Ellsberg paradox. Assume, for example, that $\phi(x,y) = 1/2\,[u(x)v(y) + v(x)u(y)]$ with u and v taking on the same values as they did in the previous subsection. Consider the following two probability distributions over states R, a red ball drawn, W, a white ball drawn, and B, a black ball drawn:

$$\mu(R) = 1/3, \mu(W) = 1/3 - \epsilon,\ \mu(B) = 1/3 + \epsilon,$$

$$\nu(R) = 1/3, \nu(W) = 1/3 + \epsilon, \nu(B) = 1/3 - \epsilon.$$

With these beliefs, the valuation of lottery A is 12,000/9 and of lottery B $12{,}000(1/9 - \epsilon^2)$. Thus, A is preferred to B. On the other hand, the evaluation of lottery C is 10,125(4/9), whereas the lottery D scores $10{,}125(4/9 - \epsilon^2)$, which shows that the representation suggested in this article can accommodate the Ellsberg paradox.

Following Sarin and Wakker (1992) we can extend this example. Suppose an urn is filled with *100* balls. There are again *20* red balls, and the number of black or white balls sums to *40*, as does the number of white or yellow balls, and the remainder is green. Clearly, the number of green balls is equal to the number of white balls and hence the number of black and green balls also sums to *40*. Let us abbreviate colors by their first letter. Figure 2 illustrates the above information about the distribution of the colors of the balls in the urn.

As Sarin and Wakker note, if a decision maker expresses a preference for betting on R over Y (i.e., $100_R0 \succ 100_Y0$), then if she is probabilistically sophisticated (i.e., her beliefs about the likelihood of an event can be represented by a unique probability measure), then

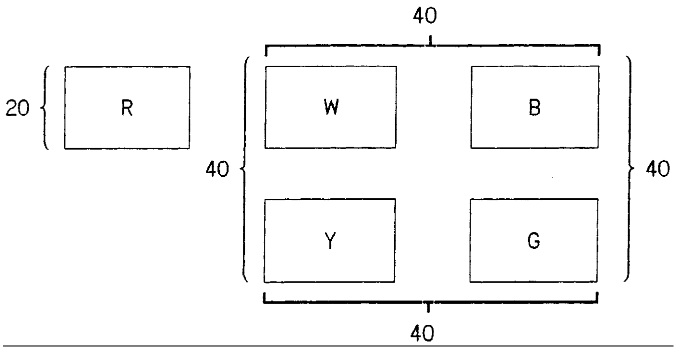

Figure 2.

she should also express preferences for betting on R∪W over Y∪W, R∪W∪B over Y∪W∪B, and so on. Sarin and Wakker argue, however, that a person may like "specificity" and dislike "ambiguity" and thus would exhibit a preference for betting: on Y∪W over R∪W, on R∪W∪B over Y∪W∪B, and on Y∪W∪B∪G over R∪W∪B∪G.

Such a preference pattern, although inconsistent with a single probability measure, can be accommodated quite naturally with our mixture symmetric representation. The following two probability measures are consistent with the data specified above:

$$\mu(R) = 0.2, \mu(Y) = \mu(G) = 0.2 + a, \mu(W) = \mu(B) = 0.2 - a,$$

$$\nu(R) = 0.2, \nu(Y) = \nu(G) = 0.2 - a, \nu(W) = \nu(B) = 0.2 + a.$$

Note that a^2 can be interpreted as a measure of the "ambiguity" the decision maker feels about the exact number of balls of each color in the urn. The representation of beliefs collapses to a single probability measure when $a = 0$. As we did above, let us assume $u(100) = v(100) = 100$, $u(0) = 10$ and $v(0) = 0$.

Using the same form for $\phi(x, y)$ as we did above, it is straightforward to show that for any a not equal to zero the decision maker would express the above pattern of betting preferences displaying an aversion to ambiguity. Let $x_E y$ denote the lottery that pays x if the event E obtains and y otherwise. For example, $100_R 0$ is the lottery that pays \$100 if the ball drawn is Red and pays nothing otherwise. So, because

$$V(100_E 0) = 9500\mu(E)\nu(E) + 500(\mu(E) + \nu(E))$$

we have:

$$V(100_R 0) > V(100_Y 0),$$

$$V(100_{R\cup W} 0) < V(100_{Y\cup W} 0),$$

$$V(100_{RuWuB}0) > V(100_{YuWuB}0) \text{ and}$$

$$V(100_{RuWuBuG}0) < V(100_{YuWuBuG}0)$$

as required.

5. Updating and resolute choice reconsidered

The examples in section 4 show that a quadratic functional for which the separable beliefs can be represented by two probabilities measures is sufficient to accommodate both Allais- and Ellsberg-type paradoxes. The last example of that section can also be utilized to demonstrate that the individual's attitude toward ambiguity that she currently faces may be dependent on the ambiguity that she has already experienced or would have experienced in unrealized events. Moreover, this exemplifies the fact that for such *resolute* choice the individual treats experienced ambiguity in exactly the same way as anticipated ambiguity.

To make these points a little more concrete, consider a situation where a ball is drawn from the urn of the second example above. Assume that the individual is allowed to bet on two colors: one of which must be Red, White, or Black; and the other that must be Yellow or Green. That is, if the ball drawn is either of the two colors that she selected, she wins $100; otherwise she wins nothing. Assuming her preferences can be represented as in that example above, we know that she prefers to bet on either Black and Green, Black and Yellow, White and Yellow, or White and Green.

For concreteness assume she chooses to bet on Black and Green. Now assume that the ball is drawn and she is told it is neither Yellow nor Green. Moreover, she is told that, if she wishes, she can change her bet on Black to either White or Red. Our dynamically consistent (resolute) agent would choose not to change her bet, because the ambiguity experienced in betting on Green complements the ambiguity that she anticipates in retaining her bet on Black.

Introspection might suggest, however, (and this has been confirmed by very casual empiricism) that, for an ambiguous-averse individual, ambiguity that lies ahead figures more prominently than ambiguity that is past or resides in unrealized events.

Although this article has demonstrated the possibility of dynamically consistent ambiguous averse non-EU (expected utility) preferences that can rationalize resolute choice, the conclusion that we draw from the above discussion is that perhaps a more natural and promising approach for modeling dynamic choice under subjective uncertainty is explicitly to account for the structure of the resolution of uncertainty through time by specifying the individual's conditional preferences at each decision node she may encounter rather than defining a preference relation over the global set of acts.

A fully specified model is beyond the scope of this article, but we think it is suggestive to conjecture that, if her beliefs were quadratic and she knew the ball's color would be revealed as it was drawn, then she would bet on Black and Green. However, if she knew she would receive the information that the ball was either Red, White, or Black or that it was Yellow or Green, and she could change her bet, then we predict she would be indifferent to betting on Yellow or Green (regardless of what she bet out of Red, White, or Black), but that she would prefer to bet on Red as this would be an unambiguous event given her updated information.

Appendix: proof of representation theorem

Every h in $\mathcal{H}$ can be viewed as a measure on $\mathcal{X} \times S$ with the restriction that for all s in $\mathcal{S}$, ${}_x h_s(x) = 1$. $\mathcal{H}$ is a mixture space and preferences on $\mathcal{H}$ satisfy continuity, prize monotonicity and mixture symmetry.

Following Chew, Epstein, and Segal (1991, p. 142), we shall denote an indifference set $I(f) := \{h \in \mathcal{H} \mid h \approx f\}$ as planar if it is convex but not equal to $\{f\}$.

Lemma A.1. There exist no planar indifference sets for $\gtrsim$.

Proof. Consider the best outcome b and the worst outcome w in $\mathcal{X}$ (i.e., for all x in $\mathcal{X}\backslash\{b, w\}$, $b > x > w$). For any p in [0,1] denote by $L(p)\colon = p \cdot \delta^x + (1 - p) \cdot \delta^y$ the lottery that yields the best outcome b with probability p and the worst outcome w with probability $(1 - p)$. Abusing notation, let $L(p)$ also denote the constant horse-race lottery that yields $L(p)$ with certainty. By continuity and prize monotonicity it follows that for any horse-race lottery h in $\mathcal{H}$, there exists a (unique) q in [0,1] such that $h \sim L(q)$. That is, every indifference set is an element of $\{I(L(p)) \mid p \in [0,1]\}$.

Suppose that for some $\bar{p}$ in (0,1), $I(L(\bar{p}))$ is planar. Then for each non-null event A in $\mathcal{E}$, there exists a t_A in (0,1) such that for all (p, q) in $[0,1]^2$, $L(p)_A\, L(q) \sim L(\bar{p})$ iff $(p - \bar{p})t_A + (q - \bar{p})(1 - t_A) = 0$.[5] Moreover, for each such event A, there exists p_A, q_A, r_A, s_A in [0,1] with $p_A > \bar{p} > q_A$ and $r_A > \bar{p} > s_A$ such that $(p_A - \bar{p})t_A + (q_A - \bar{p})(1 - t_A) = (s_A - \bar{p})t_A + (r_A - \bar{p})(1 - t_A) = 0$. Hence $L(p_A)_A\, L(q_A) \sim L(\bar{p}) \sim L(s_A)_A\, L(r_A)$. But because by construction $L(\bar{p}) = \lambda[L(p_A)_A\, L(q_A)] + (1 - \lambda)[L(s_A)_A\, L(r_A)]$, where $\lambda = (r_A - s_A)/[(p_A - q_A) + (r_A - s_A)]$ by proper uncertainty aversion $L(\bar{p}) \gtrsim L(p_A)_A\, L(q_A)$. Moreover, there exists an event B for which $L(\bar{p}) > L(p_A)_A\, L(q_A)$. This contradicts the assumption that $I(L(\bar{p}))$ is planar.

From Lemma A.1 we have that the preference relation $>$ on $\mathcal{H}$ satisfies the hypothesis of Chew, Epstein, and Segal's Theorem 5 (1991, p. 149). Hence $\gtrsim$ can be numerically represented by a proper quadratic utility function. ■

Acknowledgments

We would like to thank Peter Wakker and an anonymous referee for comments that clarified our views and helped to improve the exposition of the article substantially. Neither of them is to be held responsible for any remaining errors and obscurities.

Notes

1. See also Rabinowicz (1994).
2. Sarin and Wakker's (1993) property of sequential consistency is related to this notion.
3. After writing this article, related work by Lo (1996) came to our attention. Lo also extends the work of Chew, Epstein, and Segal (1991) to the AA framework and notes the potential of this approach to represent dynamically consistent updated preferences.
4. Epstein and Wang (1994) and Sarin and Wakker (1993) explore this approach.
5. That is, the points in Figure 1 that correspond to horse-race lotteries in $I(L(\bar{p}))$ lie on the straight line through the point $(\bar{p}, \bar{p})$ with normal vector $(t_A, 1 - t_A)$.

References

Anscombe, F. J., and R. J. Aumann. (1963). "A Definition of Subjective Probability," *Annals of Mathematical Statistics* 34, 149–205.

Chew, S. H., L. G. Epstein, and U. Segal. (1991). "Mixture Symmetry and Quadratic Utility," *Econometrica* 59, 139–164.

Ellsberg, D. (1961). "Risk, Ambiguity and the Savage Axioms," *Quarterly Journal of Economics* 75, 643–669.

Epstein, L. G., and M. le Breton. (1993). "Dynamically Consistent Beliefs Must Be Bayesian," *Journal of Economic Theory* 61, 1–22.

Epstein, L. G., and T. Wang. (1994). "Intertemporal Asset Pricing Under Knightian Uncertainty," *Econometrica* 62, 283–322.

Hammond, P. (1988). "Consequentialist Foundations for Expected Utility," *Theory and Decision* 25, 25–78.

Howard, R. A. (1992). "In Praise of the Old-Time Religion." In W. Edwards (ed.), *Utility Theories*: *Measurements and Applications*. Boston, MA: Kluwer Academic Publishers.

Karni, E., and D. Schmeidler. (1991). "Atemporal Dynamic Consistency and Expected Utility Theory," *Journal of Economic Theory* 54, 401–408.

Lo, K. C. (1996). "Weighted and Quadratic Models of Choice Under Uncertainty," *Economic Letters* 50, 381–386.

Machina, M. (1989). "Dynamic Consistency and Non-Expected Utility Models of Choice under Uncertainty," *Journal of Economic Literature* 27, 1622–1668.

McClennen, E. (1990). *Rationality and Dynamic Choice*. Cambridge, England: Cambridge University Press.

Rabinowicz, W. (1994). *To Have One's Cake and Eat It*: *How to Make Sequential Choices When One's Preferences Violate Expected Utility Axioms*. Unpublished manuscript, Department of Philosophy, Uppsala University.

Sarin, R. K., and P. P. Wakker. (1992). "A General Result for Quantifying Beliefs," *Econometrica* 62, 683–686.

Sarin, R. K., and P. P. Wakker. (1993). *Consistency in Dynamic Choice Situations with Nonexpected Utility Models*. Unpublished manuscript.

Schmeidler, D. (1989). "Subjective Probability and Expected Utility Without Additivity," *Econometrica* 57, 571–587.