

A minimal extension of Bayesian decision theory

Ken Binmore¹

Published online: 22 May 2015
© Springer Science+Business Media New York 2015

Abstract Savage denied that Bayesian decision theory applies in large worlds. This paper proposes a minimal extension of Bayesian decision theory to a large-world context that evaluates an event E by assigning it a number $\pi(E)$ that reduces to an orthodox probability for a class of measurable events. The Hurwicz criterion evaluates $\pi(E)$ as a weighted arithmetic mean of its upper and lower probabilities, which we derive from the measurable supersets and subsets of E . The ambiguity aversion reported in experiments on the Ellsberg paradox is then explained by assigning a larger weight to the lower probability of winning than to the upper probability. However, arguments are given here that would make anything but equal weights irrational when using the Hurwicz criterion. The paper continues by embedding the Hurwicz criterion in an extension of expected utility theory that we call expectant utility.

Keywords Bayesian decision theory · Expected utility · Non-expected utility · Upper and lower probability · Hurwicz criterion · Alpha-maximin

1 Preview

Bayesian decision theory was created by [Savage \(1954\)](#) in his ground-breaking *Foundations of Statistics*. He emphasizes that the theory is not intended for universal application, observing that it would be “preposterous” and “utterly ridiculous” to apply the theory outside what he calls a *small world* (Savage 1954, p. 16). It is clear that Savage would have regarded the worlds of macroeconomics and finance as large, but it is not so clear just how complex or surprising a world needs to be before he

✉ Ken Binmore
k.binmore@ucl.ac.uk

¹ Economics Department, Bristol University, 8 Woodland Road, Bristol BS8 1TB, UK

would have ceased to count it as small. Nor does he offer anything definitive on how rational decisions are to be made in a large world.¹

In recent years, a substantial literature has developed that offers various proposals on how to extend Bayesian decision theory to at least some kinds of large worlds. This paper offers another theory of the same kind that deviates more from the Bayesian orthodoxy than my book *Rational Decisions* (Binmore [2009, Chapter 9]) but remains only a minimal extension of the standard theory. It should be emphasized that the paper only discusses rational behavior, which may or may not reflect how people behave in practice.

1.1 Decision problems

A decision problem can be modeled as a function

$$D : A \times B \rightarrow C,$$

where A is a space of feasible actions, B is a space of possible states of the world whose subsets are called events, and C is a space of consequences that we assume contains a best outcome \mathcal{W} and a worst outcome \mathcal{L} . Each action is assumed to determine a finite gamble \mathbf{G} of the form

$$\mathbf{G} = \begin{array}{|c|c|c|c|c|} \hline \mathcal{P}_1 & \mathcal{P}_2 & \mathcal{P}_3 & \cdots & \mathcal{P}_m \\ \hline E_1 & E_2 & E_3 & \cdots & E_m \\ \hline \end{array} \tag{1}$$

in which $\mathcal{E} = \{E_1, E_2, \dots, E_m\}$ is a partition of the belief space B , and the prize \mathcal{P}_i is understood to result when the event E_i occurs. Different symbols \mathcal{P}_i for a prize need not represent different outcomes of a gamble. It is taken for granted that the order in which the columns of (1) appear is irrelevant, and that it does not matter whether columns with an empty E_i are inserted or deleted.

We follow the standard practice of assuming that the decision maker has a preference relation \preceq defined on whatever set G of gambles need to be considered. If \preceq satisfies sufficient rationality (or consistency) postulates, it can be described by a Von Neumann and Morgenstern (VN & M) utility function $u : G \rightarrow \mathbb{R}$. It is usually assumed that the gamble \mathbf{G} whose prizes are all \mathcal{P} can be identified with \mathcal{P} . The standard VN&M utility function $U : C \rightarrow \mathbb{R}$ of the orthodox theory can then be identified with the restriction of u to C . The assumptions of Bayesian decision theory then allow $u(\mathbf{G})$ to be expressed as the expected utility:

¹ Savage (1954) own candidate is the minimax regret criterion. Manski (2013) offers a vigorous defense of this proposal, but it seems to me a non-starter for a theory of rational decision because it fails to satisfy the Independence of Irrelevant Alternatives.

$$u(\mathbf{G}) = \sum_{i=1}^m U(\mathcal{P}_i) p(E_i), \quad (2)$$

where p is a (subjective) probability measure.

1.2 Non-expected utility

As in other generalizations of Bayesian decision theory, we replace the subjective probability $p(E)$ by a less restrictive notion that we denote by $\pi(E)$. As with $p(E)$ in the orthodox theory, $\pi(E)$ is defined as the utility $u(\mathbf{S})$ of the simple gamble \mathbf{S} that yields \mathcal{W} if the event E occurs and otherwise yields \mathcal{L} . The surrogate probability $\pi(E)$ need not be additive, but we do not call it a non-additive probability to avoid confusion with [Schmeidler \(1989, 2004\)](#) well-known theory. The utility $u(\mathbf{G})$ similarly fails to reduce to an expected utility of the form (2), but we make no use of the Choquet integral.

The version of Savage's sure-thing principle offered as Postulate 5 generalizes the linear expression (2) to a multivariate polynomial in the quantities $x_i = U(\mathcal{P}_i)$ ($i = 1, 2, \dots, m$). We refer to this extension of expected utility as *expectant utility* for the reasons given in Sect. 4.4.

Because we confine attention to minimal extensions of the Bayesian orthodoxy, it is easy to make assumptions that tip the theory back into Bayesianism. Postulate 8 is such an assumption. If it is imposed on all gambles in G , then π must be a probability and we recover the expected utility formula (2). However, Sect. 6 argues against imposing Postulate 8 in a large-world context. Instead, a class \mathcal{M} of measurable events is defined using the criterion that gambles constructed only from events in \mathcal{M} satisfy Postulate 8. The restriction of π to \mathcal{M} then satisfies the requirements of a probability, and so we denote it by p . We then say that the events in \mathcal{M} are not only measurable, but have been *measured*. Requiring that expected utility should be maximized for gambles constructed only from events in \mathcal{M} imposes restrictions on the coefficients of the multivariate polynomial representation of expectant utility (Sect. 6.1).

Ambiguity versus uncertainty There is more than one way of interpreting unmeasured events in this model. The more orthodox assumes that the decision maker would be able to assign a subjective probability to all unmeasured events if better informed, but her ignorance prevents her settling on a particular value of this probability. This *ambiguity interpretation* needs to be compared with the wider *uncertainty interpretation*, which allows that it may be intrinsically meaningless to assign probabilities to some unmeasured sets. The model of this paper is intended to be applicable even with the uncertainty interpretation.

1.3 Hurwicz criterion

The work outlined in Sect. 1.2 is prefixed by asking how a surrogate probability $\pi(E)$ might be constructed from p on the assumption that all the decision-maker's information has already been packaged in the subjective probabilities she has assigned to the measured events in \mathcal{M} . All that she can then say about an event E is that it has outer measure $\bar{p}(E)$ and inner measure $\underline{p}(E)$ (Sect. 2.2). Following [Good \(1983\)](#),

Halpern and Fagin (1992), and Suppes (1974) and others, we identify $\bar{p}(E)$ with the upper probability of E and $\underline{p}(E)$ with its lower probability.

The Hurwicz criterion (Hurwicz 1951; Chernoff 1954; Milnor 1954; Arrow and Hurwicz 1972) expresses $\pi(E)$ as a weighted arithmetic mean of $\bar{p}(E)$ and $\underline{p}(E)$:

$$\pi(E) = (1 - \alpha) \underline{p}(E) + \alpha \bar{p}(E) \quad (0 \leq \alpha \leq 1). \quad (3)$$

The ambiguity aversion reported in experiments on the Ellsberg paradox is sometimes explained by taking $\alpha < \frac{1}{2}$ in Eq. (3). Ambiguity-loving behavior then corresponds to $\alpha > \frac{1}{2}$ and ambiguity neutrality to $\alpha = \frac{1}{2}$. However, Theorem 3 suggests that, when upper and lower probabilities are identified with outer and inner measures, only $\alpha = \frac{1}{2}$ is viable.

Alpha-maximin Equation (3) assigns a utility to any simple gamble. With the ambiguity interpretation, there is a natural extension called α -maximin to the case of a general gamble \mathbf{G} . One first computes the expected utility of \mathbf{G} for all probability distributions that are not excluded by the decision-maker's information. The utility of \mathbf{G} is then taken to be a weighted arithmetic mean of the infimum and supremum of this set of expected utilities. Arguments for this conclusion are given by Ghirardato et al. (2002) and Klibanoff et al. (2005). My earlier work (Binmore 2009) assumes the same extension of the simple Hurwicz criterion to the general case, but the theory offered in Sect. 5 of the current paper is not consistent with α -maximin.

2 Unmeasured sets

The type of large-world scenario to which the theory offered in this paper is intended to apply is reviewed in Sect. 4. Earlier sections examine a particular functional form for a surrogate probability $\pi(E)$ in order to prepare the ground.

2.1 Measure and probability

A (sigma) algebra \mathcal{M} of measurable subsets of a state space B is defined to be closed under complements and countable unions. We deviate slightly from the standard definition in allowing \mathcal{M} to be empty, noting that $\mathcal{M} \neq \emptyset$ implies $\{\emptyset, B\} \subseteq \mathcal{M}$. In proving theorems, we assume that B is a finite set, so that proving something for the finite case also proves it for the countable case. We stick with the countable definitions because our results all extend to the infinite case when appropriate continuity assumptions are made.

A probability measure p on \mathcal{M} is a countably additive function $p : \mathcal{M} \rightarrow [0, 1]$ for which $p(\emptyset) = 0$ and $p(B) = 1$. When a probability measure p on \mathcal{M} has been identified, we say that the events in \mathcal{M} have been measured. We shall use \mathcal{N} to denote a larger algebra of unmeasured sets—events for which an extension of p from \mathcal{M} to \mathcal{N} may never be identified.

What kind of probability? At least three concepts of probability can be distinguished (Gillies 2000). Probabilities can be objective (long-run frequencies), subjective (in the sense of Savage), or epistemic (logical degrees of belief). All the probabilities of this

paper are either objective or subjective, on the understanding that when an objective probability is available—as with roulette wheels or dice—the decision-maker’s subjective probability coincides with its objective probability. We have nothing to say about the epistemic or logical probabilities (credences) used in attempts to solve the general problem of scientific induction.

Casino example Unmeasured sets are usually only mentioned when studying Lebesgue measure, where the emphasis is on the paradoxes that they can engender. However, there is nothing paradoxical about following example.

A blind anthropologist from Mars finds herself at a roulette table in Monte Carlo.² The only betting is on $LOW = \{1, 2, \dots, 18\}$ or $HIGH = \{19, 20, \dots, 36\}$. She hears the croupier saying various things but only his announcements of LOW or $HIGH$ seem relevant because only then does she hear the clink of chips being transferred. She therefore restricts her attention to the two events in the collection $\mathcal{S} = \{LOW, HIGH\}$ that we would regard as a knowledge partition of our state space $B = \{1, 2, \dots, 36\}$ (Binmore 2009, p. 140). Eventually, she attaches subjective probabilities to these events, which we take to be $p(LOW) = p(HIGH) = \frac{1}{2}$.

A new player now enters the casino and starts betting on $ODD = \{1, 3, \dots, 35\}$ or $EVEN = \{2, 4, \dots, 36\}$. Our Martian therefore refines her partition to $\mathcal{S}' = \{E_1, E_2, E_3, E_4\}$, where $E_1 = ODD \cap LOW$, $E_2 = ODD \cap HIGH$, $E_3 = EVEN \cap LOW$, and $E_4 = EVEN \cap HIGH$. This paper studies her decision problem at this point, before she has had the opportunity to formulate subjective probabilities for the events in her new knowledge partition.

The decision maker’s algebra of measured sets is $\mathcal{U} = \{\emptyset, LOW, HIGH, B\}$, where $LOW = E_1 \cup E_2$ and $HIGH = E_3 \cup E_4$. She also distinguishes a larger algebra \mathcal{V} containing some unmeasured sets, which consists of all unions of the elements of the partition $\{E_1, E_2, E_3, E_4\}$ of B . For example, E_1 and $E_2 \cup E_4 = EVEN$ are unmeasured. She does not recognize the algebra \mathcal{W} of all subsets of B because she is unaware that we regard the state space as B . The gambles in the set G she wishes to study are therefore only those that can be constructed from events in \mathcal{V} .

If the decision maker never needs to revise her knowledge partition again, she will not go wrong by proceeding as though her state space was simply $\{E_1, E_2, E_3, E_4\}$. We are in a similar situation when we restrict attention to the algebra \mathcal{W} of all subsets of the state space $B = \{1, 2, \dots, 36\}$. Why not take B to be the set of all physical states that determine where the roulette ball stops? Why not all possible quantum states of the universe? In brief, our choice of state space is a matter of modeling convenience.

Hausdorff’s paradox Vitali proved that some sets of points on a circle are not Lebesgue measurable.³ Lebesgue measure—which is countably additive—can be extended as a finitely additive rotation-invariant measure to all subsets of the circle. But no similar escape is available when the circle is replaced by a sphere (whose group of rotational symmetries is non-Abelian). Hausdorff showed that a sphere can

² European roulette has no 00. Bets on LOW or $HIGH$ remain on the table when 0 occurs.

³ Vitali’s argument needs the Axiom of Choice, without which all sets on the circle can be taken to be Lebesgue measurable (assuming “inaccessible cardinals” exist—Solovay 1970). But to deny the Axiom of Choice would seem to deny the ethos that led Savage to insist on the relevance of large worlds by assuming that our current formalism is adequate to describe anything that nature might throw at us.

be partitioned into three disjoint sets, A , B , and C , each of which can not only be rotated onto either of the other two, but also—paradoxically—onto the *union* of the other two (Wagon 1985).⁴ The ambiguity interpretation of Sect. 1.3 cannot then easily be sustained because a rotation-invariant extension π of Lebesgue measure would have to satisfy $\pi(A \cup B) = \pi(A) = \pi(B)$. Hausdorff’s three sets are therefore more than unmeasured—they cannot be measured in a manner consistent with Lebesgue measure.

2.2 Inner and outer measure

A minimal extension of Bayesian decision theory should only use information already packaged in the subjective probabilities assigned to measured events. Ideally, the surrogate probability of a unmeasured event E should therefore depend only on its inner measure $\underline{p}(E)$ and its outer measure $\overline{p}(E)$. The outer measure of E is the infimum of the measures $p(F)$ of all measured supersets F of E . Its inner measure is the supremum of the measures $p(F)$ of all measured subsets F of E . The inner and outer measures of a measured set are equal.

In the casino example, $\underline{p}(E_1) = 0$ and $\overline{p}(E_1) = \frac{1}{2}$. In Hausdorff’s paradox, $\underline{p}(A) = 0$ and $\overline{p}(A) = 1$ (Binmore 2009, p. 177).

It follows from their definitions that \overline{p} is subadditive and \underline{p} is superadditive. If M is measured, it is also true that $\underline{p}(M \cap E) + \overline{p}(M \cap \sim E) = \underline{p}(M)$. Thus,

$$\underline{p}(E) + \overline{p}(\sim E) = 1. \tag{4}$$

Inner measures are seldom mentioned in modern measure theory, which uses the Carathéodory criterion:

$$\overline{p}(E \cap M) + \overline{p}(E \cap \sim M) = \overline{p}(E) \tag{5}$$

for all E (measurable or not) to define the measurable sets M in an expansion \mathcal{M}^* of \mathcal{M} . For (5) to hold for our measured sets, we would need to assume that such an expansion has already been carried out. A dual treatment in terms of inner measures shows that (5) also holds with \overline{p} replaced by \underline{p} (Halmos 1950).

Lotteries A gamble in which all the events that determine the prizes are measured is called a lottery. Roulette wheels and dice are examples in which the probabilities are objective. A lottery determined by just two events H and T will be called a weighted coin. We write $p(H) = h$ and $p(T) = t$.

We allow gambles \mathbf{G} in which the prizes \mathcal{P}_i are lotteries. It is then important to remember that the assumption of Bayesian decision theory sometimes known as the replacement axiom will not be available. So it will not necessarily be true that the decision maker is indifferent between \mathbf{G} and the gamble obtained by exchanging a prize that is a lottery by another independent prize with the same Von Neumann and

⁴ His theorem has a countable set of exceptional points, but the later work of Banach and Tarski perfected his result on the way to proving their even more spectacular paradox.

Morgenstern utility. We nevertheless follow [Anscombe and Auman \(1963\)](#) in taking for granted VN&M’s theory of rational decision under risk:

Postulate 1 *Some version of the Von Neumann and Morgenstern theory of rational decision under risk applies to lotteries whose prizes are gambles in G .*

Postulate 1 implies that each gamble \mathbf{G} has a VN&M utility $u(\mathbf{G})$ that we systematically normalize so that $u(\mathcal{L}) = 0$ and $u(\mathcal{W}) = 1$.

Independence We take devices like weighted coins or roulette wheels to be the norm for lotteries by requiring that each lottery be independent of *everything* else in the model. Where it is necessary to be precise, we can tack a new lottery space L onto the state space, so that B is replaced by $L \times B$. Our strong independence requirement is then operationalized by endowing $L \times B$ with the product measure constructed from the probability measures defined on L and B . In the finite case, this is done by first constructing all “rectangles” of the form $S \times T$, where S and T are measured sets in L and B . The measured sets in $L \times B$ itself are the finite unions of all such rectangles, which can always be expressed as the union of *disjoint* rectangles. The product measure on $L \times B$ is then defined to be the sum of the products $p(S)p(T)$ for all rectangles $S \times T$ in such a union of disjoint rectangles. (The countable case is not much harder).

The reason for rehearsing this standard construction of a product measure is to clarify why it is easy to work out the inner and outer measures of sets like $(H \times E) \cup (T \times F)$ when E and F are possibly unmeasured sets in B , and H and T are the possible outcomes when a weighted coin is tossed. Because of our strong independence assumption that all lotteries are independent of all measured sets in B , we have, for example, that

$$\underline{p}(\{H \times E\} \cup \{T \times F\}) = h \underline{p}(E) + t \underline{p}(F), \tag{6}$$

$$\overline{p}(\{H \times E\} \cup \{T \times F\}) = h \overline{p}(E) + t \overline{p}(F). \tag{7}$$

3 Simple gambles

We now pursue the implications of Postulate 1 for simple gambles of the form:

$$\mathbf{S} = \begin{array}{|c|c|} \hline \mathcal{L} & \mathcal{W} \\ \hline \sim E & E \\ \hline \end{array}, \tag{8}$$

where $\sim E$ is the complement of the set E in the state space B .

With the normalization $u(\mathcal{L}) = 0$ and $u(\mathcal{W}) = 1$, we can extend a probability p given on an algebra \mathcal{M} of measured sets to a surrogate probability π defined on a larger algebra \mathcal{N} of possibly unmeasured sets by writing

$$\pi(E) = u(\mathbf{S}).$$

To exploit this definition, we need to link the assumptions of Postulate 1 (about lotteries with gambles as prizes) with the methodology for calculating inner and outer measure reviewed in Sect. 2.2. The next postulate does so by allowing a compound gamble to be treated merely as a complicated way of writing a gamble of the form (1).⁵

Postulate 2 *The procedure by means of which a gamble is implemented is irrelevant.*

It follows from Postulate 2 that the gambles

$$\mathbf{M} = \begin{array}{|c|c|} \hline \mathcal{L} & \mathcal{W} \\ \hline \sim(H \times E) & H \times E \\ \hline \end{array}; \quad \mathbf{N} = \begin{array}{|c|c|} \hline \mathcal{L} & \begin{array}{|c|c|} \hline \mathcal{L} & \mathcal{W} \\ \hline \sim E & E \\ \hline \end{array} \\ \hline \sim H & H \\ \hline \end{array} \tag{9}$$

have the same utility. Using Postulate 1 to evaluate the lottery **N**, we have that

$$\pi(H \times E) = p(H)\pi(E). \tag{10}$$

3.1 Hurwicz criterion

Suitably adapted versions of the axiom systems given by Milnor (1954) and Klibanoff et al. (2005) identify $\pi(E)$ as defined above with the Hurwicz criterion of (3). This section offers further arguments leading to the same conclusion. Theorem 3 suggests that only the ambiguity-neutral version of the Hurwicz criterion is viable for a rational decision maker. Other considerations favoring the case $\alpha = \frac{1}{2}$ are mentioned later.

Example 2 What happens in the casino example when π is given by the Hurwicz criterion? Because $\pi = p$ on \mathcal{U} , we have that $\pi(\emptyset) = 0$, $\pi(B) = 1$, and $\pi(\text{LOW}) = \pi(\text{HIGH}) = \frac{1}{2}$. For events in \mathcal{V} , $\pi(E_i) = \frac{1}{2}\alpha$, $\pi(\sim E_i) = \frac{1}{2}(1 + \alpha)$, and $\pi(F) = \alpha$ when F has two elements but is not LOW or HIGH. In the Hausdorff paradox, $\pi(A) = \pi(B) = \pi(C) = \pi(A \cup B) = \pi(B \cup C) = \pi(C \cup A) = \alpha$.

Minimal information For a minimal extension of Bayesian decision theory, $\pi(E)$ should depend only on the information packaged in $\bar{p}(E)$ and $\underline{p}(E)$. We build this requirement into the next postulate together with some regularity assumptions. In this postulate, $D = \{(x, y) : 0 \leq x \leq y \leq 1\}$.

Postulate 3 *There exists an increasing function $v : D \rightarrow \mathbb{R}$ that is homogeneous of degree 1 and continuously differentiable on D with the property that for all events E in \mathcal{N} ,*

$$\pi(E) = v(\underline{p}(E), \bar{p}(E)). \tag{11}$$

⁵ Postulate 2 does not say that the way the gamble (1) itself is structured is irrelevant, otherwise it would conflict with our later denial of Postulate 8.

A function v is homogeneous of degree 1 if $v(zx, zy) = zv(x, y)$ ($0 \leq z \leq 1$). This property is justified by (10), whose left side is $v(p(H)\underline{p}(E), p(H)\overline{p}(E))$ and whose right side is $p(H)v(\underline{p}(E), \overline{p}(E))$.

Theorem 1 *Postulates 1, 2 and 3 imply the Hurwicz criterion (3).*

Proof By Postulate 3, we can differentiate (10) partially with respect to x . Cancel a factor of z and then take the limit as $z \rightarrow 0+$ to obtain $v_1(x, y) = v_1(0, 0)$.⁶ Similarly, $v_2(x, y) = v_2(0, 0)$. Integrating these equations with the boundary condition $v(0, 0) = 0$ yields the theorem with $1 - \alpha = v_1(0, 0)$ and $\alpha = v_2(0, 0)$. \square

Postulate 3a *As Postulate 3 with v additive instead of homogeneous.*

We defend the assumption that v is additive by appealing to the Carathéodory criterion (5) and its dual, in which inner measure replaces outer measure. When these hold, any π given by the Hurwicz criterion satisfies

$$\pi(E \cap M) + \pi(E \cap \sim M) = \pi(E) \tag{12}$$

for any measured event M . Theorem 2 provides a converse.


Theorem 2 *Postulates 1, 2 and 3a imply the Hurwicz criterion (3).*

Proof Given that $v(x_1 + x_2, y_1 + y_2) = v(x_1, y_1) + v(x_2, y_2)$, we have $v(x, y) = v(x, 0) + v(0, y)$. The continuously differentiable functions $v(x, 0)$ and $v(0, y)$ are each additive, and hence⁷ $v(x, 0) = \beta x$ and $v(0, y) = \gamma y$. Appealing to the additivity of v again, $\beta = \gamma$. \square

3.2 Ambiguity neutrality

The next postulate is always true when E is a measured event and it would therefore seem inevitable with the ambiguity interpretation (which says that although we may not know how to measure a set, it nevertheless has a measure). However, the postulate implies that only the ambiguity-neutral Hurwicz criterion is viable in our setting.

Postulate 4 *Given any three independent tosses of a weighted coin in which the outcomes $H_1, H_2,$ and H_3 each occur with probability h :*



⁶ All limits in Postulate 3 are taken from within D . If $(0, 0)$ were excluded from the set D , the theorem would fail as $v(x, y) = x^{1-\alpha} y^\alpha$ would remain a possibility.

⁷ $f(x + y) = f(x) + f(y)$ implies $f'(x + y) = f'(x)$. So $f'(y) = f'(0)$ and $f(y) = f'(0)y$.

Theorem 3 *Postulates 1, 2, and 4 imply the arithmetic Hurwicz criterion with $\alpha = \frac{1}{2}$, provided that events can be found with any lower and upper probabilities.*

Proof By (4), the VN&M utility of the left-hand side of Postulate 4 is

$$\pi(\{H_1 \times \sim E\} \cup \{H_2 \times E\}) = v(h(1 - P) + hp, \quad h(1 - p)h + hP),$$

where $p = \underline{p}(E)$ and $P = \overline{p}(E)$. The utility of the right-hand side is

$$\pi(\{H_3 \times \sim E\} \cup \{H_3 \times E\}) = \pi(H_3 \times B) = \pi(H_3) = h.$$

Writing $\Delta = P - p$, it follows that

$$v(h(1 - \Delta), h(1 + \Delta)) = h.$$

The equations $x = h(1 - \Delta)$ and $y = h(1 + \Delta)$ define a bijection $\phi : (0, 1)^2 \rightarrow D^o$ provided that enough unmeasured sets E exist to ensure that each value of Δ in $(0, 1)$ is available. So for all (x, y) in the interior D^o of the domain D of v ,

$$v(x, y) = \frac{1}{2}(x + y).$$

This proof of the Hurwicz criterion does not depend on the smoothness of v , □

Example 3 Applying the Hurwicz criterion with $\alpha = \frac{1}{2}$ in the casino example, we find that π is a probability measure that extends p from \mathcal{U} to \mathcal{V} . This is an unusual result. For example, if we replace \mathcal{V} by \mathcal{W} , then $\pi(\{i\}) = \frac{1}{4}$ for all i , so that $\pi(\{1\}) + \pi(\{2\}) + \dots + \pi(\{36\}) = 9 \neq 1 = \pi(\{1, 2, \dots, 36\})$. However, (4) implies that it is always true that the Hurwicz criterion with $\alpha = \frac{1}{2}$ satisfies⁸

$$\pi(E) + \pi(\sim E) = 1. \tag{13}$$

4 Philosophy

Having offered a concrete example of a possible functional form for the surrogate probability π , it is now necessary to consider how this or another π might fit within an extension of expected utility theory. To this end, this section reviews the philosophical basis of the enterprise.

4.1 Aesop’s principle

The assumptions of most rationality theories take for granted that it would be irrational for a decision maker to allow her preferences over the prizes in C , her beliefs over the

⁸ The identity does not hold if E is unmeasured and $\alpha \neq \frac{1}{2}$, because $\pi(E) + \pi(\sim E) = 1$ implies that $1 = (1 - \alpha)p + \alpha P + (1 - \alpha)(1 - P) + \alpha(1 - p) = 1 + (2\alpha - 1)(P - p)$, where p and P are the lower and upper probabilities of E .

states in B , and her assessment of what actions are available in A to depend on each other. Aesop's story of the fox and the grapes is one of a number of fables in which he makes this point (Binmore 2009, pp. 5–9).

We follow the standard practice of assuming that enough is known about the choices that a rational decision maker would make when faced with various decision problems and that it is possible to summarize them in terms of a (revealed) preference relation \preceq defined over the set of all gambles. In doing so, it is taken for granted that the action space A currently available does not alter the decision-maker's view on whether an action a that yields the gamble \mathbf{G} would be preferred to an action b that leads to the gamble \mathbf{H} . We are then free to focus on ensuring that the decision-maker's preferences over C do not influence her beliefs over B and that her beliefs over B do not influence her preferences over C . In Bayesian decision theory, this requirement is reflected by the separation of preferences and beliefs into utilities and probabilities in the expected utility formula (2).

States of mind Economists sometimes say that people prefer umbrellas to ice-creams on rainy days but reverse this preference on sunny days. Introducing such “state-dependent preferences” is harmless in most applications, but when foundational issues are at stake, it is necessary to get round such apparent violations of Aesop's principle somehow. The approach taken here is to identify the set C with the decision-maker's *states of mind* rather than with physical objects. An umbrella in an unreformed consequence space is then replaced by the two states of mind that accompany having an umbrella-on-a-sunny-day or having an umbrella-on-a-rainy-day.

4.2 Consistency

If your betting behavior satisfies Savage (1954) axioms, his theory of subjective probability deduces that you will act as though you believe that each subset of B has a probability.

Savage's axioms are consistency requirements. Everybody would doubtless agree that rational decisions should ideally be consistent with each other, but should the ideal of consistency necessarily take precedence over everything else? Physicists, for example, know that quantum theory and relativity are inconsistent where they overlap, but they live with this inconsistency rather than abandon the accurate predictions that each theory provides in its own domain.

Savage understood that consistency is only one desideratum for a theory of rational decision. In identifying rationality with consistency, he therefore restricted the range of application of his theory to *small worlds* in which intelligent people might be able to bring their original confused judgments about the world into line with each other by modifying their beliefs when they find they are inconsistent. Luce and Raiffa (1957, p. 302) summarize Savage's views as follows:

Once confronted with inconsistencies, one should, so the argument goes, modify one's initial decisions so as to be consistent. Let us assume that this jockeying—

making snap judgments, checking up on their consistency, modifying them, again checking on consistency etc—leads ultimately to a bona fide, prior distribution.

I agree with Savage that, without going through such a process of reflective introspection, there is no particular virtue in being consistent at all. But if the world in which we are making decisions is large and complex, how could such a process be carried through successfully? (Binmore 2009, pp. 128–134)

Achieving consistency Calling decision makers rational does not automatically make them consistent. Working hard at achieving consistency is surely part of what rationality should entail. But until somebody invents an adequate theory of how this should best be done, we have to get by without any model of the process that decision makers use to convert their unformalized “gut feelings” into a consistent system of subjective probabilities.

It seems obvious that a rational decision maker would do best to consult her gut feelings when she has more evidence rather than less. For each possible future course of events, she should therefore ask herself, “What subjective probabilities would my gut come up with *after* experiencing these events?” In the likely event that these posterior probabilities turn out to be inconsistent with each other, she should then reassess her initial snap judgments until her posterior probabilities are massaged into consistency. After the massaging is over, the decision maker would then be invulnerable to surprise, because she would *already* have taken account of the impact that any future information might have on whatever internal process determines her subjective beliefs.

The end-product of such a massaging process will be a set of consistent posterior probabilities. According to Savage’s theory, their consistency implies that they can all be formally deduced from the same prior distribution using Bayes’ rule, which therefore becomes nothing more in this story than a book-keeping tool that saves the decision maker from having to remember all her massaged posterior probabilities. But the mechanical process of recovering the massaged posteriors from the prior using Bayes’ rule should not be confused with the (unmodeled) massaging process by means of which the prior was originally distilled from the unmassaged priors. Taking account of the massaging process therefore reverses the usual story. Instead of *beginning* with a prior, the decision-maker’s subjective input *ends* with a prior.

Handling surprises Shackle (1949) emphasizes that surprises—events that the decision maker has not anticipated might occur or be relevant—are unavoidable in macroeconomics. The same goes for other large worlds, notably the world of scientific endeavor. So what does a rational person do when unable to carry through the small-world process of massaging her way to consistency?

To answer this question in full is beyond the ambition of this paper. Its formalism is only relevant when a decision maker who has gradually put together a small world in which orthodox Bayesian decision theory has worked well in the past is first confronted with a surprise that leads her to pay attention to features of her environment that hitherto seemed irrelevant. It is by no means guaranteed, but as she gains more experience she may eventually create a new small world—larger than before but still small—in which Bayesian decision theory will again apply. But right now, all she has is the information packaged in her old small world and the brute fact that she has been taken by surprise by an event of whose possibility she had not previously taken account.

4.3 Knowledge

Arrow (1971), p. 45 tells us that each state in B should ideally be “a description of the world so complete that, if true and known, the consequences of every action would be known.” But as Arrow and Hurwicz (1972) explain, “How we [actually] describe the world is a matter of language, not of fact. Any description of the world can be made finer by introducing more elements to be described.” This paper follows both these prescriptions by assuming that the space B of states of the world in any decision problem $D : A \times B \rightarrow C$ is complete but not fully known to the decision maker. We also assume the same of the space C of states of the decision-maker’s mind.

How can Bayesian decision theory work if the decision maker does not know what decision problem she is facing? One possibility is that the decision-maker’s past experience has taught her what issues seem worth paying attention to in the context of her current problem. This experience will equip her with a (finite) sequence of questions to ask about the world around her and her own feelings. These questions determine a partition S of B and a partition T of C . The decision-maker’s knowledge after asking her questions then reduces to specifying in which element of each partition the actual states lie (Binnmore 2009, p. 358). As long as the sets $\{E_1, E_2, \dots, E_m\}$ and $\{\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_m\}$ in gambles (1) that arise in her decision problem are always coarsenings of the partitions S and T , it is then irrelevant whether she knows anything else about the spaces B and C .

In moving away from Bayesian decision theory, it will matter that this approach determines the partitions S and T *independently* of all decision problems that are currently envisaged. It will then not be appropriate to start a decision analysis with a simplification of the gamble \mathbf{G} of (1) that coarsens the partition \mathcal{E} by replacing E_i and E_j by $E_i \cup E_j$ when they yield the same prize. Other considerations aside, to do so is potentially to violate Aesop’s principle by allowing what happens in C to influence how B is structured. In Bayesian decision theory, it turns out that this violation does not matter (Theorem 5), but we are working outside this theory.

Similar considerations apply to the partition T of C , but a second point is more important. Our story makes a prize \mathcal{P}_i into a *set* of states of the mind rather than a deterministic object like an umbrella or an ice-cream. The decision maker will presumably not regard the states of mind in \mathcal{P}_i as being far apart or she would not have packaged them into the same element of the partition T , but there will necessarily be some ambiguity, not only about which state of mind in \mathcal{P}_i will actually be realized, but also about its possible dependence on the current state of the world. Such a potential violation of Aesop’s principle may be a source of uncertainty that needs to be incorporated somehow into the theory of how she makes decisions.

4.4 Expectant utility

We have discussed how a rational decision maker might construct a decision problem to which Bayesian decision theory applies. But what happens if her Bayesian updating is interrupted by a surprising event—something she did not anticipate when constructing her current system of consistent subjective probabilities? She might then be led to recognize that questions she did not ask previously (because they did not occur to

her or she thought them irrelevant) now need to be asked. Having asked them, her new knowledge can be described with new partitions S' and T' of B and C that are refinements of her old partitions S and T .

This paper intervenes at the stage when she has formulated the newly refined partitions, but before she has acquired enough further information to assign consistent subjective probabilities to the members of the new partition of B to supplement the subjective probabilities already established for the members of the old partition. The events she can describe using her old partition will be identified with the algebra \mathcal{M} of measured subsets of B . The events for which she needs the new partition will be unmeasured. The step to this new partition corresponds to introducing a new algebra \mathcal{N} that refines \mathcal{M} . The casino example of Sect. 2.1 illustrates how this might work with $\mathcal{M} = \mathcal{U}$ and $\mathcal{N} = \mathcal{V}$.

The next section proposes assumptions that replace expected utility in these new circumstances by a generalization that we call *expectant utility*. The terminology is intended to suggest that the notion is a temporary expedient requiring refinement as further information is received.

5 General postulates for gambles

This section offers a substitute for orthodox expected utility theory. The more interesting theorems depend on the following result of [Keeney and Raiffa \(1975\)](#).

5.1 Separating preferences

The following result applies to a consequence space C that can be factored so that $C = C_1 \times C_2$. The prizes in C then take the form $(\mathcal{P}, \mathcal{Q})$, where \mathcal{P} is a prize in C_1 and \mathcal{Q} is a prize in C_2 . We say that a preference relation \preceq on C evaluates C_1 and C_2 *separately* if and only if it is always true that

$$\begin{aligned} (\mathcal{P}, \mathcal{Q}) \prec (\mathcal{P}, \mathcal{Q}') &\text{ implies } (\mathcal{P}', \mathcal{Q}) \preceq (\mathcal{P}', \mathcal{Q}'); \\ (\mathcal{P}, \mathcal{Q}) \prec (\mathcal{P}', \mathcal{Q}) &\text{ implies } (\mathcal{P}, \mathcal{Q}') \preceq (\mathcal{P}', \mathcal{Q}'). \end{aligned}$$

If the consequence spaces C , C_1 , and C_2 are, respectively, replaced in this definition by the sets of all lotteries over these outcome spaces, the separation requirement is surprisingly strong. When \preceq can be represented by a VN&M utility function $u : C \rightarrow \mathbb{R}$, [Binmore \(2009, p. 175\)](#) obtains the multinomial expression

$$u = Au_1u_2 + Bu_1(1 - u_2) + Cu_2(1 - u_1) + D(1 - u_1)(1 - u_2), \tag{14}$$

where the functions $u_1 : C_1 \rightarrow \mathbb{R}$ and $u_2 : C_2 \rightarrow \mathbb{R}$ can be regarded as normalized VN&M utility functions on C_1 and C_2 . The constants in (14) are $A = u(\mathcal{W}_1, \mathcal{W}_2)$, $B = u(\mathcal{W}_1, \mathcal{L}_2)$, $C = u(\mathcal{L}_1, \mathcal{W}_2)$, and $D = u(\mathcal{L}_1, \mathcal{L}_2)$.

The generalization of (14) to the case when C can be factored into m components instead of 2 presents no new difficulty. For example, the expression for $m = 4$ has

sixteen terms of which a typical term is

$$u(\mathcal{W}_1, \mathcal{L}_2, \mathcal{L}_3, \mathcal{W}_4) u_1(1 - u_2)(1 - u_3)u_4.$$

The general formula for $u(\mathcal{P}_1, \dots, \mathcal{P}_m)$ is the following multinomial expression in which $x_1 = u_1(\mathcal{P}_1), x_2 = u_2(\mathcal{P}_2), \dots, x_m = u_m(\mathcal{P}_m)$:

$$\sum_{i_1=0}^1 \cdots \sum_{i_m=0}^1 u(\mathcal{X}_1^{i_1}, \dots, \mathcal{X}_m^{i_m}) y_1^{i_1} \cdots y_m^{i_m}, \tag{15}$$

where

$$\mathcal{X}_j^i = \begin{cases} \mathcal{L}_j & \text{if } i = 0 \\ \mathcal{W}_j & \text{if } i = 1 \end{cases} \quad \text{and} \quad y_j^i = \begin{cases} x_j & \text{if } i = 0 \\ 1 - x_j & \text{if } i = 1 \end{cases}.$$

Separating preferences involving gambles The preceding results on separating preferences are now applied to gambles over lotteries using the following version of Savage’s sure-thing principle.

Postulate 5 *For all gambles with lotteries as prizes:*

$$\begin{array}{|c|c|c|c|c|} \hline \mathbf{L}_1 & \cdots & \mathbf{L}_i & \cdots & \mathbf{L}_m \\ \hline E_1 & \cdots & E_i & \cdots & E_m \\ \hline \end{array} < \begin{array}{|c|c|c|c|c|} \hline \mathbf{L}_1 & \cdots & \mathbf{K}_i & \cdots & \mathbf{L}_m \\ \hline E_1 & \cdots & E_i & \cdots & E_m \\ \hline \end{array}$$

implies $\begin{array}{|c|c|c|c|c|} \hline \mathbf{K}_1 & \cdots & \mathbf{L}_i & \cdots & \mathbf{K}_m \\ \hline E_1 & \cdots & E_i & \cdots & E_m \\ \hline \end{array} \leq \begin{array}{|c|c|c|c|c|} \hline \mathbf{K}_1 & \cdots & \mathbf{K}_i & \cdots & \mathbf{K}_m \\ \hline E_1 & \cdots & E_i & \cdots & E_m \\ \hline \end{array}$

Applying Postulate 5 to any particular gamble \mathbf{G} , we find that $u(\mathbf{G})$ can be written in the form of Eq. (15), where the (normalized) utility function u_i is used to evaluate the i th prize in the gamble \mathbf{G} . The next postulate—justified by Aesop’s principle—removes the dependence of u_i on E_i . It says that it only matters what you get and not how you get it.

Postulate 6 *The utility functions u_i are the same for all non-empty events E_i in all gambles \mathbf{G} .*

We can therefore write $U(\mathcal{P}) = u_i(\mathcal{P})$ for any non-empty event E_i and so recover the VN&M utility function

$$U : C \rightarrow \mathbb{R}$$

of Sect. 1. Theorem 4 then replaces the standard expected utility formula (2).

Theorem 4 *For a fixed partition \mathcal{E} , Postulates 1, 5, and 6 imply*

$$u(\mathbf{G}) = \sum_{i_1=0}^1 \cdots \sum_{i_m=0}^1 u(\mathcal{X}_1^{i_1}, \dots, \mathcal{X}_m^{i_m}) y_1^{i_1} \cdots y_m^{i_m}, \tag{16}$$

where

$$x_j^i = \begin{cases} \mathcal{L} & \text{if } i = 0 \\ \mathcal{W} & \text{if } i = 1 \end{cases} \quad \text{and} \quad y_j^i = \begin{cases} U(\mathcal{P}_j) & \text{if } i = 0 \\ 1 - U(\mathcal{P}_j) & \text{if } i = 1 \end{cases} .$$

Theorem 4 leaves much room for maneuver in assigning values to the coefficients of Eq. (16). The next postulate is a minimal restriction.

Postulate 7 *For a fixed partition \mathcal{E} , successively replacing occurrences of \mathcal{L} by \mathcal{W} in $u(\mathcal{L}, \mathcal{L}, \dots, \mathcal{L})$ never decreases the decision-maker’s utility.*

Postulate 7 allows the decision maker to be so pessimistic that she regards any possibility of losing as equivalent to always losing, so that only the coefficient $u(\mathcal{W}, \mathcal{W}, \dots, \mathcal{W}) = 1$ in Eq. (16) is non-zero. In this case, $u(\mathbf{G}) = x_1 x_2 \dots x_m$, where $x_i = U(\mathcal{P}_i)$. Alternatively, the decision maker may be so optimistic that she regards any possibility of winning as equivalent to always winning so that all the coefficients in (16) are 1 except for $u(\mathcal{L}, \mathcal{L}, \dots, \mathcal{L}) = 0$. In this case, $u(\mathbf{G}) = 1 - (1 - x_1)(1 - x_2) \dots (1 - x_m)$.

At the other extreme, we can recover Bayesian decision theory by choosing the coefficients in (16) appropriately. In the case when all the elements of \mathcal{E} are regarded as interchangeable—as envisaged by the principle of insufficient reason—it is natural to propose that a coefficient u in (16) should be set equal to k/m , where k is the number of \mathcal{W} s in its argument. The formula then collapses to $(x_1 + x_2 + \dots + x_m)/m$, which is the expected utility of \mathbf{G} when all the elements E_i of the partition \mathcal{E} are assigned probability $1/m$. What if $u = f(k/m)$, where $f : [0, 1] \rightarrow \mathbb{R}$ is increasing and continuous? When $x_i = x$ ($i = 1, 2, \dots, m$), Eq. (16) reduces to

$$u(\mathbf{G}) = \sum_{k=0}^m f\left(\frac{k}{m}\right) \binom{m}{k} x^k (1-x)^{m-k} \rightarrow f(x) \quad \text{as } m \rightarrow \infty$$

by a theorem of [Widder \(1941\)](#), p. 152.

5.2 Reduction to surrogate probabilities

The next postulate is almost enough to convert the preceding theory into another foundation for Bayesian decision theory.

Postulate 8 *If the same prize \mathcal{P} results from multiple events in the partition that determines prizes in a gamble \mathbf{G} , then the new gamble that results from replacing these events by their union is indistinguishable from \mathbf{G} .*

The following equation is an example of how Postulate 8 works.

\mathcal{L}	\mathcal{P}	\mathcal{P}	=	\mathcal{L}	\mathcal{P}	.	(17)
$\sim(E \cup F)$	E	F		$\sim(E \cup F)$	$E \cup F$		

The wording of Postulate 8 is intended to include the requirement that when Theorem 4 is used to work out the expected utility of gambles like those of (17), then it does not matter whether the partition \mathcal{E} is taken to be $\{\sim(E \cup F), E, F\}$ or $\{\sim(E \cup F), E \cup F\}$. The strong implications are explored in Sect. 5.3.

Postulate 8 is sometimes assumed without comment, but our notion of expectant utility dispenses with it in favor of the following weaker version.⁹

Postulate 9 *If the same prize \mathcal{P} results from multiple events in the partition that determines prizes in a gamble \mathbf{G} in which all the prizes are either \mathcal{W} or \mathcal{L} , then the new gamble that results from replacing these events by their union is indistinguishable from \mathbf{G} .*

Postulate 9 is needed to bridge the gap between the surrogate probability π introduced in Sect. 3 and the more general theory being developed here. For example, in defining $\pi(E)$ as $u(\mathbf{S})$, it matters that Postulate 9 implies that

$$\mathbf{S} = \begin{array}{|c|c|} \hline \mathcal{L} & \mathcal{W} \\ \hline \sim E & E \\ \hline \end{array} = \begin{array}{|c|c|c|c|c|} \hline \mathcal{L} & \mathcal{L} & \dots & \mathcal{L} & \mathcal{W} \\ \hline E_1 & E_2 & \dots & E_{m-1} & E = E_m \\ \hline \end{array}. \tag{18}$$

Why surrogate probabilities matter Postulate 9 implies that we can express the coefficients in the formula (16) for the expectant utility of a gamble as surrogate probabilities of appropriate events. For example,

$$u(\mathcal{W}, \mathcal{L}, \mathcal{W}, \mathcal{W}, \mathcal{L}) = \pi(E_1 \cup E_3 \cup E_4). \tag{19}$$

Example 4 We know the surrogate probabilities for all events in the casino example when the Hurwicz criterion is employed, so we can work out the expectant utility of any gamble constructed from events in \mathcal{V} . In the case of the gamble \mathbf{J} in which the events LOW and HIGH yield prizes with VN&M utilities x and y ,

$$u(\mathbf{J}) = \frac{1}{2}x + \frac{1}{2}y,$$

which is the expected utility of \mathbf{J} because $p(\text{LOW}) = p(\text{HIGH}) = \frac{1}{2}$. Applying Theorem 4 with $\mathcal{E} = \{E_1, E_2, E_3, E_4\}$ to the gamble \mathbf{K} in which the events ODD and EVEN yield prizes with VN&M utilities x and y :

⁹ Postulate 9 does not even imply that a gamble can be identified with \mathcal{P} when all the prizes are \mathcal{P} . But when $\mathcal{M} \neq \emptyset$ and so B is a measured event, this conclusion follows from requiring that Bayesian decision theory holds for gambles constructed only from measured events.

$$u(\mathbf{K}) = xy(1 - 2\alpha) + \alpha x + \alpha y .$$

Unless $x = y = 0$ or $x = y = 1$, $u(\mathbf{K}) < \frac{1}{2}x + \frac{1}{2}y$ when $\alpha < \frac{1}{2}$, which can be seen as a kind of ambiguity aversion. However, only the case $\alpha = \frac{1}{2}$ satisfies the constraints on π listed in Example 6 for the formula of Theorem 4 to be compatible with the validity of Bayesian decision theory for gambles constructed only from events in $\mathcal{U} = \{\emptyset, \text{LOW}, \text{HIGH}, B\}$.

Example 5 Consider a gamble \mathbf{G} in which the three sets $A, B,$ and C of the Hausdorff paradox yield prizes with VN&M utilities $x, y,$ and z . Applying α -maximin on the assumption that all probability distributions are possible yields

$$u(\mathbf{G}) = (1 - \alpha) \min\{x, y, z\} + \alpha \max\{x, y, z\} .$$

If we use the Hurwicz criterion in Theorem 4 with $\mathcal{E} = \{A, B, C\}$, all three sets and their unions in pairs have the same surrogate probability $\pi = \alpha$. Expanding the left side of the equation $(x + 1 - x)(y + 1 - y)(z + 1 - z) = 1$, we find that

$$u(\mathbf{G}) = (1 - \alpha)xyz + \alpha\{1 - (1 - x)(1 - y)(1 - z)\} ,$$

which is a convex combination of the cases of extreme pessimism and extreme optimism mentioned in Sect. 5.1.

5.3 Expected utility

It is natural to ask under what conditions π is additive.

Postulate 10 *The range of $U: C \rightarrow \mathbb{R}$ contains at least three points.*

Theorem 5 *Postulates 1, 5, 6, 8, and 10 imply that*

$$\pi(E \cup F) = \pi(E) + \pi(F) ,$$

for all disjoint events E and F .

Proof Apply Theorem 4 with $m = 2$ to the right-hand side of (17), whose utility is therefore $x\pi(E \cup F)$, where $x = U(\mathcal{P})$. Apply Theorem 4 with $m = 3$ to the left-hand side of (17), whose utility is therefore $x^2\pi(E \cup F) + x(1 - x)\{\pi(E) + \pi(F)\}$. Postulate 8 says that these two quantities are equal, and so

$$x(1 - x)\{\pi(E \cup F) - \pi(E) - \pi(F)\} = 0 .$$

By Postulate 10, there is a value of x other than 0 or 1 so the theorem follows. □

It is now easy to show that (16) reduces to the standard expected utility formula. For example, expanding (16) in terms of $x_i = U(\mathcal{P}_i)$ when $m = 3$, the coefficient of x_1 is

$\pi(E_1)$. The coefficient of x_1x_2 is $\pi(E_1 \cup E_2) - \pi(E_1) - \pi(E_2) = 0$. The coefficient of $x_1x_2x_3$ is

$$1 - \pi(E_1 \cup E_2) - \pi(E_1 \cup E_2) - \pi(E_1 \cup E_2) + \pi(E_1) + \pi(E_2) + \pi(E_3) = 0.$$

Only the coefficients of the linear terms can therefore be non-zero. We quote the general result as a theorem:

Theorem 6 *Postulates 1, 5, 6, 8, and 10 imply that*

$$u(\mathbf{G}) = \sum_{i=1}^m \pi(E_i) U(\mathcal{P}_i).$$

Proof The proof requires looking at the utility of the gamble

$$\mathbf{G}_k = \begin{array}{|c|c|c|c|c|} \hline \mathcal{L} & \mathcal{P}_1 & \mathcal{P}_2 & \cdots & \mathcal{P}_k \\ \hline E_0 & E_1 & E_2 & \cdots & E_k \\ \hline \end{array} \tag{20}$$

in which we eventually take $E_0 = \emptyset$. By (10), $u(\mathbf{G}_1) = U(\mathcal{P}_1)\pi(E_1)$. To prove Theorem 6 by induction, it is then necessary to show that $u(\mathbf{G}_{k+1}) = u(\mathbf{G}_k) + U(\mathcal{P}_{k+1})\pi(E_{k+1})$.

Write $x_i = U(\mathcal{P}_i)$. Theorem 4 says that $u_k = u(\mathbf{G}_k)$ can be expressed as a sum of 2^k terms, each of which is a product of a coefficient $\pi(S)$ multiplied by k factors that are either x_i ($i \in I$) or $1 - x_i$ ($i \notin I$), where I runs through all subsets of $\{1, 2, \dots, k\}$. The set S is the union of all E_i with $i \in I$. For example, $S = E_1 \cup E_3$ when $I = \{1, 3\}$. Although $\pi(\emptyset) = 0$, it is useful to retain the term $\pi(\emptyset) (1 - x_1) \dots (1 - x_k)$ corresponding to $I = \emptyset$.

Next observe that

$$u(\mathbf{G}_{k+1}) = x_{k+1}v_k + (1 - x_{k+1})u_k,$$

in which v_k is the same as u_k except that each coefficient $\pi(S)$ is replaced by $\pi(S \cup E_{k+1})$. But Theorem 5 says that $\pi(S \cup E_{k+1}) = \pi(S) + \pi(E_{k+1})$. Thus

$$u(\mathbf{G}_{k+1}) = u_k + x_{k+1}\pi(E_{k+1}),$$

because u_k reduces to 1 when each coefficient $\pi(S)$ is replaced by 1. □

The proof of Theorem 6 also shows that if π is subadditive for events that arise in \mathbf{G} , then

$$u(\mathbf{G}) \leq \sum_{i=1}^m \pi(E_i) U(\mathcal{P}_i),$$

with the inequality reversed if π is superadditive.

6 Expectant utility

Maximizing expected utility is the fundamental principle of Bayesian decision theory. To escape this conclusion, we need to deny one of the postulates from which Theorem 5 follows. We deny Postulate 8. It is then necessary to fix the partition $\mathcal{E} = \{E_1, E_2, \dots, E_m\}$ of (1) and always to calculate the expected utility of a gamble in terms of this partition. Example 4 accordingly calculates the utility of the gamble in which \mathcal{P} is won if ODD occurs and \mathcal{Q} if EVEN occurs by applying Theorem 4 to \mathbf{K} below rather than to \mathbf{L} :

$$\mathbf{K} = \begin{array}{|c|c|c|c|} \hline \mathcal{P} & \mathcal{Q} & \mathcal{P} & \mathcal{Q} \\ \hline E_1 & E_2 & E_3 & E_4 \\ \hline \end{array}; \quad \mathbf{L} = \begin{array}{|c|c|} \hline \mathcal{P} & \mathcal{Q} \\ \hline E_1 \cup E_3 & E_2 \cup E_4 \\ \hline \end{array}.$$

Why deny Postulate 8? Recall from Sect. 4.2 that the decision maker has been using Bayesian decision theory in a small world when something whose possibility she failed to anticipate takes her by surprise. She is then led to ask more questions of the state of nature, with the result that her old knowledge partition \mathcal{F} is refined to a new knowledge partition \mathcal{E} . Section 4.2 also points out that there will be a similar reassessment of what counts as a prize—a reassessment she may need to review as she gains experience of the new world in which she now finds herself. The immediate point is that a modicum of uncertainty will be built into at least some of the prizes in the new set-up. When the decision maker looks at the gambles in (17), she may therefore see only two attempts at representing a whole class of possible gambles. If she is sensitive to uncertainty (either for or against), she will then have reason to deny that these representations are necessarily the same.

Although Postulate 8 is now to be denied, we continue to maintain the weaker Postulate 9. Making this exception may seem more intuitive when \mathcal{L} and \mathcal{W} are regarded as extreme states of mind outside our normal experience, but the major reason for making the assumption is that the approach being developed would otherwise depart too much from orthodox Bayesian decision theory to count as minimal. In particular, we would not be able to summarize the decision-maker’s attitude to uncertainty simply in terms of a surrogate probability π .

6.1 Defining measured sets

If Postulate 8 does not always hold, we can use it to define an algebra \mathcal{M} so that the restriction p of π to \mathcal{M} is a probability measure on \mathcal{M} . We simply look for a coarsening \mathcal{F} of \mathcal{E} for which Postulate 8 holds when the gambles considered only depend on events from \mathcal{F} . The collection of all unions of events in \mathcal{F} will then serve as the collection \mathcal{M} of measured sets.¹⁰

¹⁰ Note that we are asking more than that π be a measure on \mathcal{M} . In particular, although our standard normalization of VN&M utility functions ensures that $\pi(\emptyset) = 0$ and $\pi(B) = 1$ and so π is always a measure on $\{\emptyset, B\}$, it need not be true that $\{\emptyset, B\} \subseteq \mathcal{M}$. Example 5 is a case in which \mathcal{M} is taken to be empty (Sect. 2.1).

We thereby create a small world within our model. As in Sect. 4.2, we can imagine that this small world has been established through the decision-maker’s past experience, and that she has only just realized that she needs to consider a larger world. In this larger world, the partition \mathcal{F} has to be expanded to \mathcal{E} , and therefore, unmeasured sets now need to be taken into account.

Example 6 The requirement that Bayesian decision theory remains valid for gambles constructed only from events in the algebra generated by \mathcal{F} imposes constraints on the coefficients of the formula (16). As in (19), these coefficients are the values of $\pi(E)$ for all E in the algebra generated by \mathcal{E} . In the casino example, $\mathcal{F} = \{\text{LOW}, \text{HIGH}\}$. We take $p(\text{LOW}) = p$ and $p(\text{HIGH}) = 1 - p$, and ask that (16) with $m = 4$ reduces to expected utility for gambles in which LOW and HIGH yield prizes with respective VN&M utilities x and y . If the resulting expression is written as a polynomial in two variables, some of the coefficients must be zero (provided that enough values of x and y are available). Simplifying a little using the Carathéodory criterion (5), we obtain that the necessary constraints are

$$\begin{aligned} \pi(\sim E_1) + \pi(\sim E_2) &= 2 - p \\ \pi(\sim E_3) + \pi(\sim E_4) &= 1 + p \\ \pi(E_1 \cup E_4) + \pi(E_2 \cup E_3) &= 1. \end{aligned}$$

The Hurwicz criterion satisfies these constraints if and only if $\alpha = \frac{1}{2}$.

6.2 Refining the fundamental partition?

Without Postulate 8, it matters what is taken as the fundamental partition \mathcal{E} of the state space B . However, one can always refine \mathcal{E} to a new partition \mathcal{D} with the help of any independent lottery \mathbf{L} . Simply take \mathcal{D} to be the product of $\mathcal{E} = \{E_1, E_2, \dots, E_m\}$ and the collection $\mathcal{L} = \{L_1, L_2, \dots, L_n\}$ of all possible outcomes of \mathbf{L} . Does such a switch of the fundamental partition from \mathcal{E} to \mathcal{D} alter the expectant utility of gambles constructed only from events in \mathcal{E} ? The answer is *no* for the Hurwicz criterion with $0 \leq \alpha \leq 1$ because generalizations of (6) and (7) with more than two terms imply that

$$\pi(\{F_1 \times M_1\} \cup \dots \cup \{F_k \times M_k\}) = \pi(F_1)p(M_1) + \dots + \pi(F_k)p(M_k), \quad (21)$$

for all F_i in the algebra generated by \mathcal{E} and all (measured) M_i in the algebra generated by \mathcal{L} .

7 Conclusion

The paper proposes an extension of Bayesian decision theory that is sufficiently minimal that it allows gambles to be evaluated in terms of the VN&M utilities of the prizes and surrogate probabilities of the events that determine the prizes. The theory reduces to expected utility theory under certain conditions, but when these conditions are not met, the surrogate probabilities need not be additive, and the formula for the utility of

a gamble need not be linear. Several arguments are given for restricting attention to surrogate probabilities given by the Hurwicz criterion, with the upper and lower probabilities of an event taken to be its outer and inner measures generated by a subjective probability measure given on some subclass of all relevant events. However, only the ambiguity-neutral Hurwicz criterion satisfies all our requirements.

Acknowledgments I am grateful for funding from the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013)/ERC grant 295449. I am also grateful to David Kelsey and Karl Schlag for commenting on a previous version of the paper.

References

- Anscombe, F., & Auman, R. (1963). A definition of subjective probability. *Annals of Mathematical Statistics*, 34, 199–205.
- Arrow, K. (1971). *Essays on the theory of risk bearing*. Chicago: Markham.
- Arrow, K., & Hurwicz, L. (1972). An optimality criterion for decision-making under ignorance. In C. Carter & J. Ford (Eds.), *Uncertainty and expectations in economics*. Oxford: Basil Blackwell.
- Binmore, K. (2005). *Natural justice*. New York: Oxford University Press.
- Binmore, K. (2009). *Rational decisions*. Princeton: Princeton University Press.
- Chernoff, H. (1954). Rational selection of decision functions. *Econometrica*, 22, 422–443.
- Ghirardato, P., Maccheroni, F., Marinacci, M. (2002). Ambiguity from the differential viewpoint. Social Science Working Paper 1130.
- Gillies, D. (2000). *Philosophical theories of probability*. London: Routledge.
- Good, I. J. (1983). *Good thinking: The foundations of probability and its applications*. Minneapolis: University of Minnesota Press.
- Halmos, P. (1950). *Measure theory*. Princeton: Van Nostrand.
- Halpern, J., & Fagin, R. (1992). Two views of belief: Belief as generalized probability and belief as evidence. *Artificial Intelligence*, 54, 275–3178.
- Hurwicz, L. (1951). Optimality criteria for decision making under ignorance. Cowles Commission Discussion Paper, Statistics 370.
- Keeney, R., & Raiffa, H. (1975). Additive value functions. In J. Ponsard, et al. (Eds.), *Theorie de la decision et applications*. Paris: Fondation national pour L'Enseignement de la Gestion des Entreprises.
- Klibanoff, P., Marinacci, M., & Mukerji, S. (2005). A smooth model of decision-making under ambiguity. *Econometrica*, 73, 1849–1892.
- Luce, R., & Raiffa, H. (1957). *Games and decisions*. New York: Wiley.
- Manski, C. (2013). *Public policy in an uncertain world*. Cambridge, MA: Harvard University Press.
- Milnor, J. (1954). Games against nature. In R. Thrall, C. Coombs, & R. Davies (Eds.), *Decision processes*. New York: Wiley.
- Savage, L. (1954). *The foundations of statistics*. New York: Wiley.
- Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica*, 57, 571–585.
- Schmeidler, D. (2004). Subjective probability and expected utility without additivity. In I. Gilboa (Ed.), *Uncertainty in economic theory: Essays in honor of David Schmeidler's 65th birthday*. London: Routledge.
- Shackle, G. (1949). *Expectation in economics*. Cambridge: Cambridge University Press.
- Solovay, R. (1970). A model for set theory in which every set of real numbers is Lebesgue measurable. *Annals of Mathematics*, 92, 1–56.
- Suppes, P. (1974). The measurement of belief. *Journal of the Royal Statistical Society*, 2, 160–191.
- Wagon, S. (1985). *The Banach-Tarski paradox*. Cambridge: Cambridge University Press.
- Widder, D. (1941). *The Laplace transform*. Princeton: Princeton University Press.