

EXPECTED UTILITY WITHOUT UTILITY

ABSTRACT. This paper advances an interpretation of Von Neumann–Morgenstern’s expected utility model for preferences over lotteries which does not require the notion of a cardinal utility over prizes and can be phrased entirely in the language of probability. According to it, the expected utility of a lottery can be read as the probability that this lottery outperforms another given independent lottery. The implications of this interpretation for some topics and models in decision theory are considered.

KEY WORDS: Expected utility, cardinal utility, benchmark, risk attitude, stochastic dominance.

1. INTRODUCTION (AND DISAPPEARANCE)

This paper advances an alternative interpretation of the Von Neumann and Morgenstern’s expected utility model for preferences over lotteries. In short, we show that the expected utility model need not be based on the notion of a cardinal utility function over prizes and can in fact be entirely phrased in the language of probability. From here, we move to examine what the proposed interpretation has to say or to add about some important topics in expected utility theory and about some non-expected utility models.

We begin with the introduction of this new interpretation and the (simultaneous) disappearance of the von Neumann–Morgenstern (NM) utility function. Given a compact interval $C \subset \mathbf{R}$ (with nonempty interior) of monetary outcomes, a random variable X taking values in C is called a (monetary) *lottery* on C . We write $\Delta(C)$ to denote the set of monetary lotteries on C and $X \sim F$ to indicate that X has cumulative distribution function F . Note that by definition the probabilities of a lottery are exogenously given; when this is not the case, we will speak of *acts*.

Consider the problem of representing the preference relation \succeq of an agent named Nemo over the set $\Delta(C)$ of monetary lotteries. If \succeq

is complete, reflexive, transitive, and continuous (with respect to the topology of the weak convergence), it can be shown that there exists a continuous real-valued Bernoulli index $B : \Delta(C) \rightarrow \mathbf{R}$ such that $X \succeq Y$ if and only if $B(X) \geq B(Y)$, for any X, Y in $\Delta(C)$. See Theorem 1 in Grandmont (1972).

When \succeq satisfies also the axiom of independence, we obtain the expected utility model for preferences over lotteries. According to this, there exists an NM-utility function $U : C \rightarrow \mathbf{R}$ (unique up to positive affine transformations) such that the Bernoulli index representing \succeq can be written as

$$B(X) = \int_C U(x)dF(x)$$

so that Nemo ranks two monetary lotteries $X \sim F$ and $Y \sim G$ according to the rule

$$(1) \quad X \succeq Y \text{ if and only if } \int_C U(x)dF(x) \geq \int_C U(x)dG(x)$$

which compares the expected utilities of X and Y .

We assume as usual that the NM-utility function U over C is nonconstant and increasing, and furthermore that U is (at least right) continuous. As we recall in the Appendix, this latter assumption can be derived axiomatically by continuity of \succeq with respect to the topology of the weak convergence. Under this set of assumptions, U is bounded on C and by an appropriate positive affine transformation we can normalize its range to be the interval $[0, 1]$. Henceforth, U turns out to be a (nonconstant) increasing and continuous function such that $U(\inf C) = 0$ and $U(\sup C) = 1$ or, more simply, a *bona fide* cumulative distribution function on C .

As it is well known, this implies that on some appropriate probability space there exists a random variable $V \sim U$, which can always be taken to be independent of any lottery in $\Delta(C)$. See Theorem 14.1 in Billingsley (1986). It follows that the Bernoulli index used to rank lotteries from $\Delta(C)$ can be written

$$(2) \quad \begin{aligned} B_V(X) &= \int_C U(x)dF(x) = \int_C \Pr(V \leq x)dF(x) \\ &= \Pr(X \geq V) \end{aligned}$$

and (1) can be restated as

$$(3) \quad X \succeq Y \text{ if and only if } \Pr(X \geq V) \geq \Pr(Y \geq V)$$

so that two lotteries X, Y in $\Delta(C)$ are compared with respect to their likelihood of yielding outcomes which are no worse than those generated by the stochastically independent lottery V . In other words, we can interpret the expected utility model as a procedure that ranks lotteries on the basis of their probability to outperform an (independent) stochastic benchmark.

We call V an *NM-benchmark* and its cumulative distribution function $U(x) = \Pr(V \leq x)$ an *NM-distribution*. Moreover, to provide our interpretation with a name, note that in (2) we are averaging probabilities rather than utilities. Hence, we will speak of the (*expected*) *probability* model for preferences over lotteries instead of the expected utility model.

The two models coincide under the assumption that the NM-benchmark V is stochastically independent of the lotteries to be ranked. In this case, in fact, a quick glance to (2) should make it obvious that the linear functional representing Nemo's preferences is the same for both models. Mathematically, thus, stochastic independence of the benchmark implies that which of the two interpretations is chosen makes no difference.

On the other hand, the (*expected*) probability model has higher generality because, if we drop this assumption, (3) defines a Bernoulli index of the nonexpected utility type. We discuss this possibility in Section 7 but, unless explicitly stated, we maintain the assumption of stochastic independence elsewhere in the paper. We hope to show, in fact, that replacing the NM-utility function with the NM-distribution can lead to new insights even when the two models are mathematically equivalent.

The rest of the paper assumes the (*expected*) probability viewpoint and develops as follows. Section 2 deals with the interpretation of a NM-benchmark. Section 3 considers risk aversion and asset integration. Section 4 studies the uniqueness properties of the NM-distribution U . Section 5 characterizes constant risk attitude by means of noninformed benchmarks. Section 6 develops a unifying definition for both univariate and multivariate first-order stochastic dominance. Section 7, finally, looks at some nonexpected utility models.

2. INTERPRETATIONS OF THE BENCHMARK

A crucial component of the (expected) probability model is the NM-benchmark V . In this section we offer three complementary interpretations for its meaning. Our point of view is not to unveil how V is obtained but what it could represent. In particular, we hope to suggest which kind of process may lead to *construct* preferences which satisfy the (expected) probability model for preferences over lotteries. To be fair, it should be said that none of these interpretations seem completely convincing to us; yet, they seem to be more illuminating than the intuitions about the cardinality of the NM-utility function U offered by the expected utility model.

The first interpretation is that V represents the (random) value of *opting out* of a decision problem involving lotteries. The idea is the following. Presumably, Nemo does not go around with a ready-made package of preferences over lotteries to use whenever it is fit to do so. Rather, he explores (and maybe constructs) his preferences whenever this is necessary to make a choice. Hence, the problem of representing Nemo's preferences over lotteries presupposes the existence of a problem of choice among lotteries. Maybe Nemo is a sophomore who has agreed to participate in some economic experiment; maybe he is a gambler playing roulette in Las Vegas; or maybe he is just working out *Gedanken* experiments to improve his understanding of decision theory. But he *is* participating or playing or thinking. He is in the decision.

If Nemo is in, he might as well stay out. The sophomore could spend the four hours that the experiment will last working part time; the gambler could go out of the casino and play in the stock market; the researcher could stop working on expected utility and write a paper on some ongoing hot topic. Any of these acts opts out of the decision problem and leads onto a different path where the choice is among acts and not about lotteries. We can interpret V as the subjective probabilistic assessment describing what Nemo believes would happen if his choice among lotteries would be turned into a choice about acts. Moreover, the stochastic independence of V would be a consequence of the plausible assumption that lotteries are independent of the acts involved in opting out. Note also that opting out cannot be reduced to a lottery and therefore the objection

that the option of staying out should be included in Nemo's decision problem does not apply.

The second interpretation is that Nemo would actually like to use a 'probability-of-ruin' criterion, which ranks lotteries according to the probability that they outperform some deterministic threshold θ in C . However, when this threshold is only imperfectly known, Nemo makes a probabilistic assessment about its distribution U and considers the expected value of the criterion over all possible thresholds. Formally,

$$E [\Pr(X \geq \theta)] = \int_C \Pr(X \geq \theta) dU(\theta) = \Pr(X \geq V)$$

In this case, the NM-distribution U represents Nemo's uncertainty about the 'correct' threshold to apply.

In the third interpretation, V stands for some standard of reference which defines which prizes is 'fair' for Nemo to expect from the lotteries he plays or can play. More precisely, $U(x)$ represents Nemo's probability to receive a prize not greater than x in a world 'fair' to his situation. Here, the idea is that most agents have access only to a small subset of $\Delta(C)$ which depends on who they are. For instance, a very poor person can play lotto or bingo but may not play in casinos with high entrance fees; almost everybody can buy insurance policies but only a few ones can sell them. Therefore, the kind of lotteries Nemo plays depends on who he is and, in turn, this shapes his way to assess them.

According to this interpretation, Nemo learns or estimates a plausible distribution for the prizes he can earn through lotteries (given who he is and his personal history) and use this as a benchmark to rank the lotteries which he can actually play. Thus, the distribution of V may be correlated with the lotteries that Nemo plays because he can use the information about the kind of lotteries to which his current situation gives him access to infer (or revise) a plausible distribution for V . As we discuss in Section 7, an interesting consequence of this is that some well-known non-expected utility models can be recast in the framework of the (expected) probability model.

Moreover, if winning at lotto makes Nemo a millionaire and opens up new opportunities for him, he might revise drastically his assessment about V . This suggests that the NM-distribution may also be a function of Nemo's current situation (see Robson, 1992) and that

expected utility preferences over multi-stage lotteries may change over time as the outcomes of the lotteries in the initial stages become known. For lack of space, these possibilities are not considered in this paper.

3. RISK AVERSION AND ASSET INTEGRATION

According to Kahneman and Tversky (1979), most applications of the expected utility model for preferences over lotteries are based on the three tenets of expectation, asset integration, and risk aversion. Expectation states that Nemo's preferences over lotteries are represented by their expected utility. Asset integration says that the domain of the NM-utility function are Nemo's final asset positions (including current wealth) rather than his gains or losses. Risk aversion states that Nemo prefers the degenerate lottery $\delta_{E(X)}$ to any nondegenerate lottery X with expected value $E(X)$ and is equivalent to concavity of the NM-utility function.

Their paper (and many others) have offered several empirical refutations for each of these principles, and it is not our intention to rescue them here. However, we believe that the (expected) probability model can be used to shed some light on both risk aversion and asset integration and on how the empirical evidence bears on them.

Let us begin from risk aversion. The two most common characterizations of this property are either behavioral (Nemo prefers $\delta_{E(X)}$ to X) or mathematical (U is concave). Neither of them, however, seems to offer any explanation about why Nemo's preferences should be risk averse. The (expected) probability model provides a simple and direct answer. When the NM-distribution U is concave, Nemo attaches more probability to the occurrence of low outcomes for the NM-benchmark V . Therefore, Nemo is risk averse because he adopts a pessimistic benchmark where *'the worst outcomes are the most likely to happen'*. To put it more colorfully, Nemo is risk averse because he believes in Murphy's Law.

The principle of asset integration says that gains or losses are incorporated in the current asset position before evaluating the NM-utility function. However, since Markowitz (1952) formulated the assumption that NM-utility functions shift horizontally as a function of current wealth, several studies have produced ample evidence

that in most cases people tend to compute NM-utility functions over changes in wealth rather than on final asset positions. When this is the case, the function U is usually S-shaped: it is concave over gains but convex (and steeper) over losses, contradicting risk aversion. On the other hand, when Nemo is led to frame a decision problems in terms of final asset positions, U usually turns out to be concave.

In the expected utility interpretation, this evidence is usually reported as another proof of the well-known fact that Nemo's preferences (and hence their representation) may depend on how the choice problem is framed. However, why the framing affects Nemo's preferences remains unexplained. The (expected) probability interpretation can offer a plausible explanation.

To construct his preferences, Nemo must recall relevant information and assess a benchmark. The framing directs Nemo's attention either to his final asset positions or to changes from his current wealth w . In the first case, since his current position is irrelevant to the final position, Nemo's benchmark is likely to be independent of w . In the second case, instead, changes in current wealth relate naturally to w and Nemo might assess NM-distributions depending on his current position.

In this latter case, Nemo might plausibly view w as a focal value for his benchmark and assess for it a unimodal distribution U whose mode is (very close to) w . This would give the S-shaped distribution reported in the literature and described above. Furthermore, the greater steepness of the convex part of the distribution would follow from the assumption that Nemo's current wealth w is sufficiently low: as Nemo cannot expect to face losses much larger than his current wealth w , a NM-distribution centered around w has to be positively skewed. Since most experiments have been run using students (who typically belong to the low-wealth range of the population), this explanation implies that the higher steepness of U over losses may be due to an experimental bias.

On the other hand, consider the situation where Nemo is led to frame his decision problem in terms of final asset positions. Since he cannot rely on information about his current wealth, he might assess a benchmark based on his perception of the distribution of wealth in the population. Although this is a matter for empirical investigation,

casual evidence suggests that in this case most people would assess a distribution which is concave over most of its domain.

4. ARBITRARY PRIZES

The (expected) probability model can be extended without much difficulty to the case of an arbitrary prize set \mathcal{C} . For any prize x in \mathcal{C} , denote by δ_x the degenerate lottery on \mathcal{C} and by $\Delta(\mathcal{C})$ the set of all lotteries on \mathcal{C} . Suppose that \mathcal{C} is a separable metric space completely preordered by a preference relation \succeq_1 and that $\Delta(\mathcal{C})$ is completely preordered by \succeq_2 . As detailed in the Appendix (under slightly more general assumptions), there exists a bounded and continuous function $U : \mathcal{C} \rightarrow \mathbb{R}$ such that \succeq_2 can be represented by the Bernoulli index

$$(4) \quad B(X) = \int_{\mathcal{C}} U(x) dF(x)$$

for any lottery $X \rightsquigarrow F$ on \mathcal{C} if and only if: (i) \succeq_2 is continuous in the topology of the weak convergence; (ii) satisfies the independence axiom; and (iii) is consistent with \succeq_1 in the sense that $x \succeq_1 y$ if and only if $\delta_x \succeq_2 \delta_y$.

This latter assumption is especially important to the (expected) probability model because it implies that the NM-utility function U represents both \succeq_1 on \mathcal{C} and \succeq_2 on $\Delta(\mathcal{C})$. Hence, U is an increasing function or, more precisely, an order-homomorphism from (\mathcal{C}, \succeq_1) to (\mathbb{R}, \geq) . Barring the case when \succeq_1 is in fact an indifference relation and U is thus constant, then, $U : \mathcal{C} \rightarrow \mathbb{R}$ is a nonconstant, increasing, bounded and continuous function which can be normalized to a *bona fide* cumulative distribution function for some benchmark V on \mathcal{C} and the Bernoulli index defined in (4) is equivalent to $\Pr(X \succeq_1 V)$. This establishes the applicability of the (expected) probability model.

We can use this setting to shed light on an important puzzle associated with the expected utility model. Consider the usual statement that any strictly increasing transformation of the NM-utility function U represents Nemo's preferences on \mathcal{C} , but U is defined only up to positive affine transformations in the expected utility model. This says that the utility function U is an ordinal representation of \succeq_1 on \mathcal{C} and a cardinal representation of \succeq_2 on the subspace (equivalent

to C) of the degenerate lotteries in $\Delta(C)$. Why is it that U has this twofold nature?

According to the (expected) probability model, the answer is simple. The NM-distribution U supports a cardinal measure of an ordinal assessment. Given a lottery X and a benchmark V , the assessment of X is ordinal because the event $\{\omega : X(\omega) \succeq_1 V(\omega)\}$ involves only ordinal comparisons for each ω . But its measure is cardinal because it is formulated as the probability of this event, so that the ranking of X is given by $\Pr(X \succeq_1 V)$ and $X \succeq_2 Y$ if and only if $\Pr(X \succeq_1 V) \geq \Pr(Y \succeq_1 V)$. Let us examine this point in detail.

First, consider the ordinal assessment. In the (expected) probability model, preferences are independent of any order-preserving transformation of the domain of lotteries in the following sense. Given an order-homomorphism g from (\mathcal{C}, \succeq_1) to (\mathbb{R}, \geq) , the lottery $X \rightsquigarrow F$ on \mathcal{C} defines another lottery $g(X)$ on $g(\mathcal{C})$. Accordingly, a NM-benchmark V on \mathcal{C} induces the NM-benchmark $g(V)$ on $g(\mathcal{C})$. Moreover, for any x, y in \mathcal{C} , $x \succeq_1 y$ if and only if $g(x) \geq g(y)$ so that, for any ω ,

$$X(\omega) \succeq_1 V(\omega) \text{ if and only if } g[X(\omega)] \geq g[V(\omega)]$$

Hence, $\Pr(X \succeq_1 V) = \Pr[g(X) \geq g(V)]$ and $X \succeq_2 V$ if and only if $g(X)$ is at least as preferred as $g(V)$.

Therefore, when g is a real-valued (ordinal) utility function on \mathcal{C} representing \succeq_1 , all qualitative comparisons can be made directly on the space $g(\mathcal{C})$ of the ordinal ‘utilities’ of the prizes in \mathcal{C} . It is for this reason that in the following (except in Section 6) we assume without loss of generality that the set of prizes is a (not necessarily compact) subset $C = g(\mathcal{C})$ of monetary outcomes in \mathbb{R} .

Second, consider the cardinal measure. According to a common rendition, a cardinal utility function says that we can mark a zero point and a unit of measurement on a ‘utils yardstick’ and then use this yardstick to go around and measure the utility of different prizes in such a way that differences in utilities are meaningful. In the (expected) probability model, if x and y are in \mathcal{C} , the difference between their utilities is $U(y) - U(x) = \Pr(x \prec_1 V \preceq_1 y)$: the yardstick we are using is a ‘probability rod’ where differences in utilities measure strength of beliefs. Once we agree that zero and one mark respectively the null and the sure event, for any space which

is order-homomorphic to \mathcal{C} there is only one possible yardstick or, equivalently, a unique probability distribution.

5. CONSTANT RISK ATTITUDE

Constant absolute risk attitude designates preferences over lotteries that are independent of a change in the size of the initial wealth position. Analogously, constant relative risk attitude denotes preferences that are independent of a proportional change in the initial wealth and in the scale of the prizes.

In the expected utility framework, constant (absolute or relative) risk attitude has two main characterizations. The first one is very useful for modelling purposes and states that preferences with constant absolute risk attitude corresponds to linear or exponential NM-utility functions on \mathbb{R} , while constant relative risk attitude yields power (including linear) or logarithmic NM-utility functions on \mathbb{R}^+ . The second characterization relies on the notion of a constant coefficient of (absolute or relative) risk aversion and lacks a clear decision theoretic interpretation.

Since the expected utility and the (expected) probability models are mathematically equivalent, these two characterizations still hold for the (expected) probability interpretation. However, this latter can be used to provide two additional characterizations.

The first one that we consider brings out clearly the central role of the probability assessment underlying the stochastic benchmark $V \rightsquigarrow U$. Recall from Section 2 the interpretation of the (expected) probability model as a stochastic ‘probability-of-ruin’ criterion. Suppose that Nemo is a Bayesian who is trying to make a probabilistic assessment about the imperfectly known threshold θ , which we interpret for convenience as his minimum acceptable level of wealth.

Since the NM-distribution U over θ is independent of the lotteries that Nemo is called to play, both the sequence of lotteries faced and their realizations in the past do not provide information about θ . In the absence of other types of information about θ , then, Nemo might wish to use what Bayesian statisticians call a *noninformative prior* for θ and we call a *noninformed benchmark* V . A few different general methods have been proposed to derive noninformative probability

distributions. See Berger (1985) or Bernardo and Smith (1994) and references therein. A largely used method is based on a principle of invariance and, as it turns out, suffices to characterize the NM-distributions which exhibit constant risk attitude.

More precisely, constant absolute risk attitude and constant relative risk attitude correspond respectively to the so-called relatively location invariant and relatively scale invariant model for noninformative priors, which are well-known in Bayesian statistics. For a proof of this and other results in this section, see Castagnoli and Li Calzi (1993) which provides a more detailed discussion of the relevance of the (expected) probability model to constant risk attitude.

THEOREM 1. *A noninformed NM-benchmark exhibits constant absolute risk attitude on any compact interval C of prizes if and only if it is relatively location invariant.*

Thus, according to the (expected) probability interpretation, Nemo's preferences exhibit constant absolute risk attitude when his assessment of the benchmark is neutral with respect to information about its location. A similar result holds for constant relative risk attitude with respect to information about its scale.

THEOREM 2. *A noninformed NM-benchmark exhibits constant relative risk attitude on any compact interval C of strictly positive prizes if and only if it is relatively scale invariant.*

Naturally, the relatively location invariant NM-distributions are linear or exponential on \mathbb{R} and the relatively scale invariant NM-distributions are power (including linear) or logarithmic on \mathbb{R}^+ . Therefore, any of these NM-distributions (if nonconstant) is unbounded on its domain, raising a justified concern about the applicability of the (expected) probability model. In fact, many noninformative priors are unbounded and Bayesians have long learned to think of them as convenient shortcuts for more complicated statements.

To illustrate what we mean, let us go back to the expected utility model. Assume that U is an unbounded NM-utility function. Taking for granted that U must be increasing, it can be unbounded only if its domain C is not bounded. In this case, there exist lotteries with infinite expected utility and the preferences represented by

the expected utility model cannot be complete, contradicting one of our assumptions. See for instance Arrow (1974). A statement based on unbounded NM-utility functions, then, is either an implicit assumption of incomplete preferences or a shortcut for something else.

In the (expected) probability interpretation, the shortcut is the following. If U is defined over a non-compact set (say, \mathbb{R}) and is unbounded, U may be an improper cumulative distribution with infinite mass of the kind used in the Bayesian literature to model noninformative priors. However, since any realistic set C of prizes can be assumed compact, the NM-utility function U_C on C given by

$$U_C(x) = \begin{cases} 0 & \text{if } x < \min C \\ \frac{U(x) - U(\min C)}{U(\max C) - U(\min C)} & \text{if } \min C \leq x \leq \max C \\ 1 & \text{if } x > \max C \end{cases}$$

is bounded. For obvious reasons, we call U_C a *conditional NM-distribution*. It is clear that although a NM-distribution U can be unbounded, this is never the case for a conditional NM-distribution on a compact domain. In this language, the statement that the only NM-utility functions exhibiting constant absolute risk attitude are the (unbounded) exponential or linear functions becomes the statement that the only NM-distributions satisfying constant absolute risk attitude are characterized by (bounded) exponential or linear conditional NM-distributions U_C on any compact set C of prizes.

The second characterization that we consider is the following. Let V be a given NM-benchmark. We are interested in measuring how the Bernoulli index $B_V(X)$ for an available lottery X varies when we assess its value with respect to the alternative NM-benchmark $W = V + a$, where $a \in \mathbb{R}$. Natural candidates for this purpose are the *ratio index* $RI(X, a) = B_V(X)/B_{V+a}(X)$ and the *difference index* $DI(X, a) = B_V(X) - B_{V+a}(X)$. Note that the ratio index can take values in the extended reals.

In general, either index depends on the lottery X , on the value of a and on the compact domain C which defines the conditional NM-distribution U_C . The two classes of NM-distributions which exhibit constant absolute risk attitude are respectively obtained when one or

the other index is independent of X ; i.e., when the difference or the ratio between an assessment using V or $V + a$ does not depend on the lottery to be evaluated.

THEOREM 3. *Let V be a NM-benchmark with NM-distribution $U : \mathbb{R} \rightarrow \mathbb{R}$. The following holds:*

A1 *U is exponential if and only if $RI(X, a)$ is independent of X on any compact interval C of prizes;*

A2 *U is linear if and only if $DI(X, a)$ is independent of X on any compact interval C of prizes.*

A similar construction leads to the characterization of constant relative risk attitude. Given a NM-benchmark V , however, we are now interested in measuring how the Bernoulli index $B_V(X)$ for a feasible lottery X varies when we assess its value with respect to the alternative NM-benchmark $W = bV$, where $b > 0$. We use again the ratio index $RI(X, b) = B_V(X)/B_{bV}(X)$ and the difference index $DI(X, b) = B_V(X) - B_{bV}(X)$.

THEOREM 4. *Let V be a NM-benchmark with NM-distribution $U : (0, +\infty) \rightarrow \mathbb{R}$. The following holds:*

A3 *U is power if and only if $RI(X, b)$ is independent of X on any compact interval C of strictly positive prizes;*

A4 *U is logarithmic if and only if $DI(X, b)$ is independent of X on any compact interval C of strictly positive prizes.*

6. STOCHASTIC DOMINANCE

The (expected) probability model has an interesting application to the study of (first order) stochastic dominance for multivariate lotteries. We assume in the following that \mathbb{R}^n is endowed with the natural componentwise partial order \geq and that $C \subset \mathbb{R}^n$ is the cartesian product of compact intervals (with nonempty interior). In particular, if $n = 1$, this reduces to the same setting discussed so far. As above, we denote by $\Delta(C)$ the set of all lotteries on C .

In the univariate case with $n = 1$, the Bernoulli index associated by the (expected) probability model to a lottery X in $\Delta(C)$ is

$$B_V(X) = \Pr(X \geq V) = 1 - \Pr(X < V)$$

In the multivariate case with $n > 1$, instead, $\Pr(X \geq V)$ and $\Pr(X < V)$ need not have unit sum. This leads to consider an upper Bernoulli index

$$B_V^u(X) = \Pr(X \geq V)$$

and a lower Bernoulli index

$$B_V^l(X) = 1 - \Pr(X < V) = \Pr(X \not< V)$$

Note that the quantity $B_V^l(X) - B_V^u(X)$ is always not negative because it represents the probability that X and V are not comparable; i.e., that X outperforms V over at least one and at most $n - 1$ components of \mathbb{R}^n but is outperformed over the other components.

The two indices define respectively an upper and a lower preference relation \succeq^z ($z = u, l$) on $\Delta(C)$ described by $X \succeq^z Y$ if and only if $B_V^z(X) \geq B_V^z(Y)$ for all X, Y in $\Delta(C)$. As it is obvious, these two preference relations need not coincide in general for $n > 1$. In fact, the upper preference relation ranks X higher than Y if X has an higher probability than Y to outperform the benchmark over all the components of \mathbb{R}^n while the lower preference relation ranks X higher than Y if it has a higher probability of not being outperformed by the benchmark over all the components of \mathbb{R}^n .

Finally, note that the two indices can also be thought as the expected value respectively of the upper NM-utility function $U^u(x) = \Pr(V \leq x)$ and the lower NM-utility function $U^l(x) = 1 - \Pr(V > x)$. Moreover, the upper NM-utility function U^u is in fact a cumulative distribution function while the lower NM-utility function U^l is the complement to one of a survival function.

Let us now examine how this relates to the subject of multivariate stochastic dominance. Consider first the univariate case, where C is a subset of \mathbb{R} . As it is easy to show (see also Theorem 17.A.1 in Marshall and Olkin, 1979), given lotteries $X \rightsquigarrow F$ and $Y \rightsquigarrow G$, the following five conditions are equivalent and each of them could be taken as the definition of stochastic dominance:

D1 $\Pr(X \geq t) \geq \Pr(Y \geq t)$ for any t in C ;

D2 $\Pr(X < t) \leq \Pr(Y < t)$ for any t in C ;

D3 $EU(X) \geq EU(Y)$ for any increasing NM-utility function $U : C \rightarrow \mathbb{R}$;

D4 $g(X)$ first order stochastically dominates $g(Y)$ for any increasing function $g : C \rightarrow \mathbb{R}$;

D5 $\Pr(X \in B) \geq \Pr(Y \in B)$ for any subset B of C with increasing indicator function.

Consider now the multivariate case, when C is a subset of \mathbb{R}^n for $n > 1$. By analogy, it would seem natural to interpret ‘increasing’ in the conditions above as ‘componentwise increasing’ and obtain five characterizations of multivariate stochastic dominance. However, it is obvious that the two conditions D1 and D2 cannot be equivalent. Moreover, any of D1 and D2 is implied by (but does not imply) D3, D4, and D5, which are all equivalent. See Marshall and Olkin (1979). This missing implication is somewhat crucial: condition D3, which is based on expected utilities, does not follow from either D1 or D2, which are based on distribution functions.

Theorem 17.C.3 in Marshall and Olkin (1979) shows where is the root of the problem.

THEOREM 5. *Condition D1 holds on \mathbb{R}^n if and only if $EU(X) \geq EU(Y)$ for any distribution function U on \mathbb{R}^n . Analogously, condition D2 holds on \mathbb{R}^n if and only if $EU(X) \geq EU(Y)$ for any function U on \mathbb{R}^n which is the complement to one of a survival function.*

Hence, we can restore the equivalence $D1 \Leftrightarrow D3 \Leftrightarrow D4 \Leftrightarrow D5$ between the characterizations of multivariate stochastic dominance by interpreting ‘increasing’ as ‘increasing as a distribution function’ (or, equivalently, ‘increasing and supermodular’). Similarly, the equivalence $D2 \Leftrightarrow D3 \Leftrightarrow D4 \Leftrightarrow D5$ is restated if we interpret ‘increasing’ as ‘increasing as the complement to one of a survival function on C ’ (or, equivalently, ‘increasing and submodular’). In both cases, the relevance of the probabilistic interpretation should be clear.

Moreover, this suggests a possible way to unify the definition of stochastic dominance for the univariate and the multivariate case.

DEFINITION 6. Let X and Y be two random variables on \mathbb{R}^n . We say that X stochastically dominates Y in the upper sense if $\Pr(X \geq V) \geq \Pr(Y \geq V)$ for any random variable V stochastically independent of X and Y .

Similarly, we say that X stochastically dominates Y in the lower sense if $\Pr(X < V) \leq \Pr(Y < V)$ for any random variable V stochastically independent of X and Y .

It is clear that conditions D1 and D2 are respectively equivalent to stochastic dominance in the upper and in the lower sense. Hence, since D1 and D2 are equivalent in the univariate case, for $n = 1$ the two types of stochastic dominance coincide. Moreover, using Theorem 5, it is a simple exercise to prove the following characterization which holds for any $n \geq 1$.

THEOREM 7. *Let X and Y be two random variables on \mathbb{R}^n , with $n \geq 1$. Then X stochastically dominates Y in the upper (resp., lower) sense if and only if any of D1, D3, D4, D5 (resp., D2, D3, D4, D5) above holds with ‘increasing’ replaced by ‘increasing as a distribution function’ (resp., ‘increasing as the complement to one of a survival function’).*

7. CORRELATED BENCHMARKS

A natural generalization of the (expected) probability model is to drop the assumption of stochastic independence of the NM-benchmark. This can be done at various degrees of generality. A fairly general assumption is to allow for pairwise stochastic dependence between the benchmark $V \rightsquigarrow U$ and each lottery $X \rightsquigarrow F$. Thus, the Bernoulli index becomes

$$\begin{aligned}
 (5) \quad B_V(X) &= \Pr(X \geq V) = \int_C \Pr(V \leq x \mid X = x) dF(x) \\
 &= \int_C U(x; X) dF(x)
 \end{aligned}$$

and the NM-distribution $U(x; X)$ may depend on the lottery X to be evaluated. This more general type of Bernoulli index can be interpreted as if Nemo is using the information about the prizes offered

by lottery X to evaluate the appropriate benchmark against which X should be assessed. More precisely, the choice of the benchmark may be the outcome of a process in which Nemo combines some (prior) information and the evidence associated with X .

For practical purposes, however, this model is unmanageable and it is necessary to make some additional assumption. One possibility is to make the NM-distribution depend on the distribution function F rather than on the random variable X so that (5) becomes

$$(6) \quad B_V(X) = \Pr(X \geq V) = \int_C U(x; F) dF(x)$$

This mathematical formulation leads to a model that has recently surfaced inside the expected utility framework, where it has been explored at various degrees of generality by Becker and Sarin (1987), Viscusi (1989), Bordley and Hazen (1991), and Bordley (1992). For instance, the lottery-dependent expected utility model in Becker and Sarin (1987) hypothesizes that the dependence can be fully captured by a single parameter $\alpha_F \in \mathbf{R}$ so that $U(x; F) = U(x; \alpha_F)$.

In a spirit similar to Section 3, these studies have shown that instances of (6) are not only potentially consistent with a large part of the experimental evidence, but predict most of the observed behavior. However, they have left unclear the mechanism by which the lottery X to be evaluated might affect the NM-utility function.

The (expected) probability model, suggesting a probabilistic mechanism, may offer an appropriate language to unify these and other kinds of nonexpected utility models. Although this point calls for further research, we illustrate it with an example based on the generalized expected utility model by Machina (1982). The choice of this important general model is motivated by our opinion that its reinterpretation according to the (expected) probability model provides the intuitive appeal that lacks to its traditional presentation in the expected utility framework.

Let us go back to (6). Instead of specifying exactly how $U(\cdot; F)$ depends on F , we might be content with some qualitative assumption. In particular, since the assessment of the benchmark depends both on F and on other (prior) information, a plausible requirement is that $U(\cdot; F)$ 'changes less' than F for small changes in F . To formalize this statement, suppose that C is a bounded subset of \mathbf{R} and endow the set $\Delta(C)$ with the L^1 -norm. As argued in Machina (1982),

this norm is a natural choice because it induces the topology of weak convergence on the linear space generated by $\Delta(C)$.

For any two lotteries $X \rightsquigarrow F$ and $Y \rightsquigarrow G$ on C , then, we assume that

$$(7) \quad \|U(\cdot; F) - U(\cdot; G)\| = o(\|F - G\|)$$

In other words, for any lottery $Y \rightsquigarrow G$ in an appropriate neighborhood of $X \rightsquigarrow F$, the distributions of the benchmarks for X and Y are closer than the distributions of X and Y . Whatever process Nemo is using to revise his benchmark, it has the property that it is locally more stable than the change in lotteries as measured by the L^1 -norm.

When (7) holds, for Y in some appropriate neighborhood of X , the difference between the Bernoulli indices for lotteries X and Y can be written

$$\begin{aligned} B(Y) - B(X) &= \int_C U(x; G) dG(x) - \int_C U(x; F) dF(x) \\ &= \int_C U(x; F) [dG(x) - dF(x)] \\ &\quad + \int_C [U(x; G) - U(x; F)] dG(x) \\ &\leq \int_C U(x; F) [dG(x) - dF(x)] \\ &\quad + \|U(\cdot; F) - U(\cdot; G)\| \\ &= \int_C U(x; F) [dG(x) - dF(x)] + o(\|F - G\|) \end{aligned}$$

This gives exactly the key mathematical expression obtained in Machina (1982, p. 293–294) from the assumption that the Bernoulli index is Fréchet differentiable with respect to the L^1 -norm on the linear space generated by $\Delta(C)$. Note that this approach is not only more intuitive but yields naturally the assumption of Fréchet differentiability for the preference functional.

The point of view based on the (expected) probability model can be used to reinterpret most of the results presented in Machina (1982). For instance, in the expected utility interpretation, Theorem 2 in Machina (1982) states that Nemo always prefers a lottery X to a lottery Y which differs by a mean preserving increase in risk if and only if the NM-distribution $U(x; F)$ associated with any

lottery $X \rightsquigarrow F$ on C is concave in x . For the (expected) probability model, this holds if and only if all benchmarks associated with any lottery share the pessimistic belief that the worst outcomes are the most likely to happen.

Without going into detailed restatements for all the results in Machina (1982), we examine only the interpretation of its Hypothesis II. In the language of the expected utility model, this assumption states that the coefficient of absolute risk aversion for $U(\cdot; F)$ is greater than the coefficient of absolute risk aversion for $U(\cdot; G)$ whenever $X \rightsquigarrow F$ first-order stochastically dominates $Y \rightsquigarrow G$. As discussed in Machina (1983), the addition of Hypothesis II to the generalized expected utility model can explain some of the most prominent systematic violations of the independence axiom, like the common consequence effect, the common ratio effect, the oversensitivity to changes in small probabilities or outlying events and the utility evaluation effect. It is unfortunate that such a crucial assumption should sound so complicated.

Following the (expected) probability interpretation, instead, Hypothesis II states that when X first order stochastically dominates Y the benchmark for X is more pessimistic than the benchmark for Y . In other words, Nemo's assessment about the benchmark tends to be revised downward as the lottery evaluated becomes more appealing. This attitude reminds of the saying '*hope for the best, prepare for the worst*' and seems much closer to an intuitive explanation of Nemo's behavior.

APPENDIX

A. AN AXIOMATIZATION

We recall here the axiomatization of the expected utility model for preferences over lotteries by Grandmont (1972), which provides a set of necessary and sufficient conditions for the NM-utility function to be bounded and continuous. See also Foldes (1972).

We begin with some notation. Let \mathcal{C} be a separable metric space and $B(\mathcal{C})$ its Borel σ -field. Let $M(\mathcal{C})$ the space of all (countably additive) probability measures on the measurable space $(\mathcal{C}, B(\mathcal{C}))$, endowed with the topology of the weak convergence. For any $x \in \mathcal{C}$,

denote by δ_x the degenerate lottery on x and by D the (topological) subspace of all degenerate lotteries in $M(\mathcal{C})$.

We are now ready to state the axiomatization, which is given as Theorem 3 in Grandmont (1972).

THEOREM 8. *Let \mathcal{C} be a separable metric space, completely pre-ordered by \succeq_1 . Let Π be a closed and convex (topological) subspace of $M(\mathcal{C})$ containing D and completely preordered by \succeq_2 . A set of necessary and sufficient conditions for the existence of a bounded and continuous function $U : \mathcal{C} \rightarrow \mathbb{R}$ order-preserving with respect to \succeq_1 and such that the Bernoulli index $B : \Pi \rightarrow \mathbb{R}$ defined by $B(\pi) = \int_{\mathcal{C}} U(x) d\pi(x)$ is order-preserving with respect to \succeq_2 is:*

(consistency) for any x and y in \mathcal{C} , $x \succeq_1 y$ if and only if $\delta_x \succeq_2 \delta_y$;

(continuity) for any π_0 in Π , the sets $\{\pi \in \Pi : \pi \succeq_2 \pi_0\}$ and $\{\pi \in \Pi : \pi_0 \succeq_2 \pi\}$ are closed;

(independence) for any $\pi_1, \pi_2, \pi_3 \in \Pi$ and any real number $\alpha \in [0, 1]$, $\pi_1 \sim_2 \pi_2$ implies $\alpha\pi_1 + (1 - \alpha)\pi_3 \sim_2 \alpha\pi_2 + (1 - \alpha)\pi_3$.

ACKNOWLEDGEMENTS

We wish to thank L. Peccati, G.A. Rossi and especially an anonymous referee for their helpful comments. Financial support by M.U.R. S.T. and Bocconi University is gratefully acknowledged.

REFERENCES

- Arrow, K.J.: 1974, 'The use of unbounded utility functions in expected utility maximization: Response', *Quarterly Journal of Economics* **88**, 136–138.
- Becker, J.L. and Sarin, R.K.: 1987, 'Lottery dependent utility', *Management Science* **33**, 1367–1382.
- Berger, J.O.: 1985, *Statistical Decision Theory and Bayesian Analysis*, second edition, Springer, New York.
- Bernardo, J.M. and Smith, A.F.M.: 1994, *Bayesian Theory*, Wiley, New York and Chichester.
- Billingsley, P.: 1986, *Probability and Measure*, second edition, Wiley, New York and Chichester.
- Bordley, R.F.: 1992, 'An intransitive expectations-based Bayesian variant of prospect theory', *Journal of Risk and Uncertainty* **5**, 127–144.

- Bordley, R.F. and Hazen, G.: 1992, 'Nonlinear utility models arising from unmodelled small worlds intercorrelations,' *Management Science* **38**, 1010–1017.
- Castagnoli, E.: 1990, 'Qualche riflessione sull'utilità attesa,' *Ratio Mathematica* **1**, 51–59.
- Castagnoli, E. and Li Calzi, M.: 1993, 'Expected utility without utility: constant risk attitude', *Rendiconti del Comitato per gli Studi Economici* **30–31**, 145–160.
- Foldes, L.: 1972, 'Expected utility and continuity', *Review of Economic Studies* **39**, 407–421.
- Grandmont, J.-M.: 1972, 'Continuity properties of a von Neumann–Morgenstern utility', *Journal of Economic Theory* **4**, 45–57.
- Kahneman, D. and Tversky, A.: 1979, 'Prospect theory: an analysis of decision under risk', *Econometrica* **47**, 263–291.
- Machina, M.J.: 1982, 'Expected utility' analysis without the independence axiom', *Econometrica* **50**, 277–323.
- Machina, M.J.: 1983, 'Generalized expected utility analysis and the nature of observed violations of the independence axiom', in B.P. Stigum and F. Wenstøp (eds.), *Foundations of Utility and Risk Theory*, Reidel, Dordrecht, 263–293.
- Markowitz, H.: 1952, 'The utility of wealth', *Journal of Political Economy* **60**, 151–158.
- Marshall, A.W. and Olkin, I.: 1979, *Inequalities: Theory of Majorization and Its Applications*, Academic Press, Orlando (Florida).
- Robson, A.J.: 1992, 'Status, the distribution of wealth, private and social attitudes to risk', *Econometrica* **60**, 837–857.
- Viscusi, W.K.: 1989, 'Prospective reference theory: toward an explanation of the paradoxes', *Journal of Risk and Uncertainty* **2**, 235–264.

E. CASTAGNOLI

Istituto di Metodi Quantitativi
Università 'L. Bocconi'
20136 Milano, Italy.

M. LI CALZI

Dipartimento di Matematica Applicata
Università di Venezia
30123 Venezia, Italy.