

Stochastic dominance and opaque sweetening

Ralf M. Bader

Merton College, University of Oxford

ABSTRACT: This paper addresses the problem of opaque sweetening and argues that one should use stochastic dominance in comparing lotteries even when dealing with incomplete orderings that allow for non-comparable outcomes.

I The puzzle

Unlike other axioms of decision theory, such as the transitivity axiom, which have a plausible claim to being consistency norms that one is rationally required to comply with, the completeness axiom has been considered to be dubious since the very inception of axiomatic decision theory (cf. von Neumann and Morgenstern: 1944, p. 630). It has been adopted primarily for reasons of convenience, since it allows for a real-valued representation rather than requiring a more cumbersome representation, such as a vectorial representation (cf. von Neumann and Morgenstern: 1944, p. 29) or a set of functions approach (cf. Aumann: 1962).

There are many reasons why orderings need not be complete (whether one is concerned with an agent's preference ordering or with an axiological betterness ordering). Completeness will fail whenever one is dealing with non-comparability. Two alternatives A and B are non-comparable in terms of betterness ($A \not\bowtie B$) iff it is neither the case that $A > B$, nor $A = B$, nor $A < B$. Non-comparability can arise either because no positive value relations obtain between A and B , e.g. due to incommensurable values, or because of a fourth non-transitive value relation such as parity or imprecise equality (cf. Chang: 2002).¹ Non-comparability allows for insensitivity to mild sweetening: if A and B are non-comparable, then a slightly improved version A^+ can likewise be non-comparable with B , such that $A \not\bowtie B$ and $A^+ \not\bowtie B$ yet $A^+ > A$.²

¹When comparability is understood in terms of determinate betterness, then non-comparability can also be due to indeterminate betterness of the type discussed by Broome: 1997.

²Local forms of non-comparability (that are best thought of in terms of intersection quasi-orderings that integrate different dimensions of evaluation) only allow for insensitivity to mild sweetening. By contrast, global forms of non-comparability (according to which comparability

Insensitivity to mild sweetening gives rise to a puzzle in probabilistic cases, namely the puzzle of opaque sweetening (cf. Hare: 2010). Consider a lottery L_1 with a 50%-50% chance of A and B occurring, as well as a lottery L_2 over A^+ and B^+ , also with a 50%-50% chance. Importantly, the two lotteries are not independent but are defined for the same states of nature. The outcome that results in state of nature N_1 given L_1 is A, whereas B^+ would have eventuated in this state under L_2 , and similarly for the other outcomes. For instance, an agent is presented with two boxes. Non-comparable goods A and B are distributed according to a flip of a fair coin and \$10 are placed in addition into the right box.³ The agent does not know the outcome of the coin flip and, accordingly, assigns equal probabilities to the two possible states N_1 : A is in the left box and B^+ (i.e. B together with the \$10) is in the right box, and N_2 : B is in the left box and A^+ (i.e. A together with the \$10) is in the right box. The agent has two options. He can either opt for L_1 : pick the left box, or L_2 : pick the right box. This is a case of opaque sweetening, since it involves mild sweetening combined with opacity insofar as one is uncertain whether one ends up with A^+ or B^+ as a result of taking the sweetened option.

	N_1 (50%)	N_2 (50%)
L_1	A	B
L_2	B^+	A^+

The puzzle is the following. L_2 seems to be better than L_1 and hence the only rationally permissible option. Yet, no matter which state of nature ends up obtaining, the outcome that will be realised given L_2 will not be comparable to the alternative that would have resulted under L_1 . For instance, if N_2 obtains, then L_2 yields A^+ which is not comparable to the alternative B that would have resulted had L_1 been chosen instead. This means that L_2 will be better than L_1 despite the fact that the outcomes it yields are not better than those of L_1 in any state of nature. This leads to the puzzling result that we seem to have betterness at the level of the lotteries without betterness at the level of the outcomes. Although the outcomes are not comparable, the lotteries would appear to be comparable. This combination of ex ante betterness together with ex post non-comparability is rather puzzling.

2 Stochastic dominance

We are pulled towards two conflicting claims about the relation between the lotteries. On the one hand, L_2 appears to be better than L_1 , yet, on the other,

is an equivalence relation) allow for stronger forms of insensitivity. The arguments in this paper apply equally to local and global forms of non-comparability.

³Examples of supposedly non-comparable goods that are often cited include different careers as well as works of art. So one might, for instance, consider the contents of the boxes as being two paintings: A = Monet's 'La Promenade' and B = Manet's 'Le Printemps'.

there is also reason to think that these lotteries are not comparable. Whereas Hare: 2010 ultimately opts for the betterness claim, Schoenfield: 2014 as well as Bales, Cohen, and Handfield: 2014 put forward constraints that support the non-comparability claim. This paper argues in favour of the betterness claim.

The non-comparability verdict is due to state-by-state reasoning. One compares the outcome of L_2 in state N_i with the outcome of L_1 in that very same state. This kind of state-by-state comparison reveals that every state is such that the relevant outcomes are non-comparable. This is then taken to support the claim that the lotteries are likewise non-comparable. On this approach it matters in which particular state an outcome is realised. This means that it is not enough to be told that there is a probability P that outcome O will be realised. Rather, one needs to be told that O will be realised in state of nature N (with probability P). What speaks in favour of a particular lottery is that it brings about certain outcomes in particular states of nature, where the identity of the state of nature is of importance. Reasons for choosing a lottery are essentially tied to particular states, i.e. reasons are individuated in such a way that they concern outcomes in states of nature. This means that different states imply different reasons, even when they have the same probability of resulting in the same outcome.

This approach is unduly restrictive. If there is no evaluative interaction between outcomes and states, i.e. if the value of an outcome is not affected by the identity of the state, then it does not matter in which particular state an outcome occurs. The identity of the state is then evaluatively irrelevant and all that matters is the probability of the state. Reasons for choosing a lottery should, accordingly, be individuated in terms of prospects, i.e. in terms of outcomes with probabilities. This means that sameness of outcome and probability implies sameness of reason, even if the states differ. As a result, states having exactly the same probability can be treated interchangeably and hence can be permuted without affecting the evaluation of the lottery. Lotteries that are permutations of each other will then be equally good (where L and L' are permutations iff for every outcome X , $P(O = X|L) = P(O = X|L')$, i.e. they have the same probability profile for the same outcomes). This means that L_1 is equally good as its permutation L'_1 , which is state-by-state dominated by L_2 .

	N ₁ (50%)	N ₂ (50%)
L ₁	A	B
L' ₁	B	A

Given the transitivity of betterness (in particular given PI-transitivity), $L_2 > L'_1$ together with $L_1 = L'_1$ implies that $L_2 > L_1$. The betterness claim, accordingly, is supported on the basis of a comparison of prospects that abstracts from the identity of the states of nature and considers only their respective probabilities. L_2 is considered to be better than L_1 on the grounds that for every state N_i there is a corresponding state N_j with the same probability such that the outcome of

L_2 in N_i is better than the outcome of L_1 in N_j . This kind of stochastic reasoning is not restricted to state-by-state comparisons but allows for comparisons across equiprobable states.

On this approach one can straightforwardly use (first-order) stochastic dominance to order lotteries even when faced with non-comparable outcomes. A lottery L stochastically dominates L' if the former is such that for any outcome X the probability of receiving something that is at least as good as X is at least as great as under the latter and is strictly greater for some outcome Y .

$$\forall X, P(O \geq X|L) \geq P(O \geq X|L') \wedge \exists Y, P(O \geq Y|L) > P(O \geq Y|L')$$

If L stochastically dominates L' , then there is a probability-preserving permutation σ of a sufficiently fine-grained partition into states of nature, such that L state-by-state dominates L' relative to the permutation, i.e. L stochastically dominates L' iff for every state N_i , $L(N_i) \geq L'(\sigma(N_i))$ and for some N_j , $L(N_j) > L'(\sigma(N_j))$.

In the case at hand, L_2 stochastically dominates L_1 since both lotteries have a probability = 0.5 of yielding something that is at least as good as A and likewise for an outcome at least as good as B , yet L_2 has a probability = 0.5 of something at least as good as A^+ whereas the corresponding probability for L_1 is = 0, and again likewise for an outcome at least as good as B^+ .

The fact that there is no dominance of outcomes under any state is unproblematic. This is because taking L_2 guarantees that one ends up in a situation that is better than an equiprobable alternative in which one could have ended up. In this way, one switches the focus to how things *could* have turned out, as opposed to only focusing on how they *would* have turned out. Rather than considering what would have occurred under the same state of nature, one is concerned with what could just as well (i.e. with the same probability) have occurred. (This effectively amounts to decoupling the lotteries.) The outcome $L_2(N_1)$, though not being better than the outcome that would have resulted in the same state of nature had one opted for L_1 instead, i.e. $L_1(N_1)$, is nevertheless better than a situation, namely the outcome of L_1 given N_2 , that one could just as easily have ended up in, where the qualification ‘just as easily’ indicates that the requirement that the states are equiprobable is satisfied.

In other words, L_2 is better not because there is some state N_i such that the outcome of L_2 under N_i is better than the alternative that would have resulted under N_i given L_1 . Instead, it is better because in each state the outcome is better than an alternative that could just as easily have resulted, i.e. there is a bijection between the outcomes of L_2 and L_1 that preserves the probabilities and with respect to which L_2 dominates.

It is worth noting that stochastic dominance requires the outcomes that are comparable across the two lotteries to occur in states of nature having exactly the same probabilities. If the probabilities only differ slightly, then stochastic

dominance will not apply. For example, if the coin is not perfectly fair but slightly biased so that the probability of N_1' is 50.1% and that of N_2' is 49.9%, then it is no longer possible to establish the ordering $L_2 > L_1$ by means of stochastic dominance reasoning. Although the probability of ending up with an outcome at least as good as B/B^+ is higher in the case of L_2 , the probability of an outcome at least as good as A is higher on L_1 , thereby undermining stochastic dominance.⁴

2.1 Independent lotteries

The state-by-state approach is unduly restrictive (given that there is no evaluative interaction between the outcome and the state of nature). This can be brought out further by considering independent lotteries. When comparing actions that belong to different decision problems, it is not appropriate to insist on state-by-state reasoning. For instance, if we ask whether it is better to have the option L_1 which yields A and B under N_1 and N_2 respectively, or to have the option L_2' which yields A^+ and B^+ under different states of nature N_3 and N_4 respectively, then we are evaluating independent lotteries that are not defined for the same states of nature.

$$\begin{array}{c} \frac{N_1 (50\%) \quad N_2 (50\%)}{L_1 \quad A \quad B} \\ \\ \frac{N_3 (50\%) \quad N_4 (50\%)}{L_2' \quad A^+ \quad B^+} \end{array}$$

L_2' is intuitively better than L_1 . Yet, these lotteries are non-comparable according to state-by-state reasoning, according to which one lottery cannot be better than another unless there is some state of nature in which it results in a better outcome. This is because the relevant outcome for comparison is not defined when requiring the state of nature to be held fixed. For instance, $L_2'(N_1)$ is not defined. This implies that there is no outcome under L_2' that is comparable to $L_1(N_1)$. As a result, these lotteries will be deemed to be non-comparable when, intuitively, they can be ordered.

⁴In order to avoid a non-comparability verdict in those cases, one will have to invoke different resources. For instance, one can apply dominance reasoning to the probabilistically discounted value vector.

$$\begin{array}{l} L_1 : \langle 0.501 \times A, 0.499 \times B \rangle \\ L_2 : \langle 0.499 \times A^+, 0.501 \times B^+ \rangle \end{array}$$

$L_2 > L_1$ as long as the difference in value between A^+ and A is sufficiently large to compensate for the slightly lower probability of ending up with the former rather than the latter. This is the case if $0.499 \times A^+ \geq 0.501 \times A$. L_2 will then be at least as good as L_1 with respect to \mathcal{A} and will be strictly better with respect to \mathcal{B} . Non-comparability then only arises in cases in which the difference in probability is too great relative to the evaluative difference between A^+ and A .

Since what would have come about otherwise is not defined in the case of independent lotteries, one can only evaluate an outcome relative to what could have come about (with a certain probability). In order to compare these lotteries one needs to compare outcomes not under the same state of nature but outcomes with the same probabilities. This means that comparability is achieved by the fact that the respective states that can occur under L_1 and L'_2 have the same probabilities.

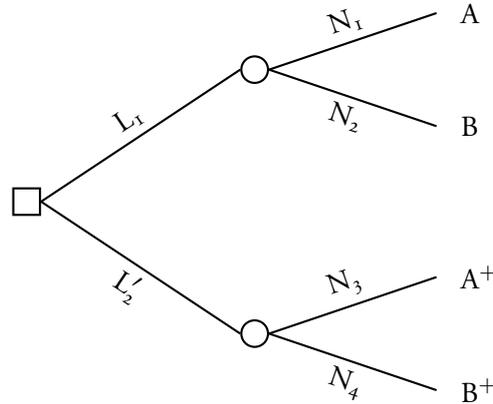
In response it can be objected that state-by-state reasoning is concerned with the rational permissibility of choices and, as such, is only meant to apply to situations in which an agent is choosing amongst a range of possible alternatives. Since decision problems do not involve independent lotteries, they are not relevant when assessing principles of rational permissibility.

Whilst traditional approaches require lotteries to share the same states of nature, given that they presuppose a universal state space that constitutes the domain for all the actions that are open to the agent, as well as an act-independent probability function over this state space, they have repeatedly been criticised for doing so and alternative models of decision-making have been proposed that do not involve these commitments. Most notably, Jeffrey-Bolker decision theory eschews these commitments. However, there are also act-dependent models that operate within an act-state framework and that do not involve such a radical departure from traditional decision theory.⁵

The theory developed by Balch and Fishburn, for instance, allows for act-dependence (cf. Balch and Fishburn: 1974). This theory starts with a set of acts and associates with each act a set of mutually exclusive and exhaustive states that are conditioned by that act. This allows us to model an agent who is faced with a choice at time t that determines the nature of subsequent chance nodes. In such cases it is contingent on the prior action at t which mutually exhaustive and exclusive states will constitute the state space at t' . Whilst the state space is act-dependent, which state eventuates is not under the control of the agent.

For example, an agent is faced with two alternative courses of actions: he can either opt for L_1 in which case a fair coin will be flipped and the agent will receive A if the coin lands heads and B if it lands tails, or for L'_2 in which case a fair six-sided die will be rolled and the agent receives A^+ if it lands 1-3 and B^+ if it lands 4-6. This means that either N_1 or N_2 will obtain if he opts for L_1 , whereas one of two different states, namely N_3 or N_4 , will obtain if he opts for L'_2 . We thus have two act-dependent state spaces: $N_{L_1} = \{N_1, N_2\}$ and $N_{L'_2} = \{N_3, N_4\}$.

⁵Whereas Jeffrey-Bolker theory is an evidential decision theory, the accounts considered below are non-standard versions of causal decision theory.



In this case, we do not have a universal state space, but different state spaces associated with different actions. Since these independent lotteries do not share states of nature but involve disjoint state spaces, it is not possible to use state-by-state reasoning. The agent instead has to make use of stochastic reasoning, which allows him to order these lotteries and rank L_2' above L_1 .

Similarly, Luce and Krantz's conditional expected utility theory treats acts as partial functions that are only defined for sub-domains of the universal state space. They thereby try to model situations in which "the decision maker controls to some extent the environment in which the consequences occur as well as the probability distribution of their occurrence" (Luce and Krantz: 1971, p. 254; also cf. Krantz et al.: 1971, ch. 8). The consequence for a conditional action L^E is not defined for a state N_i that is not in the event E on which L is conditional. For instance, $E = \{N_1, N_2\}$ whereas $F = \{N_3, N_4\}$ in the case of the coin flip and the roll of the die, such that $L_1^E(N_3)$ is not defined since $N_3 \notin E$.

	N_1	N_2	N_3	N_4
L_1^E	A	B	-	-
$L_2'^F$	-	-	A ⁺	B ⁺

Even though there is a universal state space, one cannot make use of state-by-state reasoning to compare actions that are defined for disjoint sub-domains. Given that acts are only partial functions, comparisons are not defined for states that are not shared by the events on which the acts are conditioned. Since $E \cap F = \emptyset$, L_1^E and $L_2'^F$ are non-comparable. Moreover, the state space $N = \{N_1, \dots, N_4\}$ violates the probabilistic independence condition that is presupposed by state-by-state dominance reasoning (cf. section 2.3), e.g. $P(N_3|L_1^E) = 0\%$ yet $P(N_3|L_2'^F) = 50\%$. Stochastic reasoning, however, is applicable.

In response it can be argued that these decision problems can be recast in the traditional framework in which states are act-independent. Both Balch and Fishburn as well as Luce and Krantz accept that their conditional representations can be transformed into unconditional representations by identifying a suitable underlying universal state space (cf. Fishburn: 1981, pp. 143-144 & Krantz et al.:

1971, pp. 414-416). (However, they argue that the traditional characterisation of these kinds of decision problems is significantly less natural and often not tractable due to involving very large numbers of states.)

In the case of the coin-die decision problem, for example, what are taken to be states by the conditional approach can be considered to be consequences relative to an underlying state space $S = S_1 \dots S_4$, i.e. heads = $L_1(S_1 \cup S_2)$, tails = $L_1(S_3 \cup S_4)$, die landing 1-3 = $L'_2(S_1 \cup S_3)$, and die landing 4-6 = $L'_2(S_2 \cup S_4)$.

	S_1 (25%)	S_2 (25%)	S_3 (25%)	S_4 (25%)
L_1	A	A	B	B
L'_2	A ⁺	B ⁺	A ⁺	B ⁺

Understood in this way, L_1 and L'_2 are defined for the same states. Whilst we do not have state-by-state dominance as traditionally conceived, the weakened version of dominance (advocated by Bales, Cohen and Handfield) that is specified not in terms of ‘being at least as good as’ but in terms of ‘not being worse than’ is applicable and implies $L'_2 > L_1$.⁶ This weakened dominance principle, however, leads to troublesome intransitivities (as well as failures of expansion and contraction consistency).

	N'	N''
L_A	A ⁺	B
L_B	C	B ⁺
L_C	A	C ⁺

The only options that are comparable are A and A⁺, B and B⁺, as well as C and C⁺. A pairwise comparison of the lotteries on the basis of the weakened dominance principle leads to a cyclic ordering, i.e. $L_A > L_C$, $L_C > L_B$, $L_B > L_A$.

By contrast, stochastic dominance reasoning establishes that $L'_2 > L_1$ without generating any intransitivities and without requiring any cumbersome reformulation into a representation over shared states. Moreover, if one is willing to accept the (rather controversial) claim that agents can introduce indeterminacy into the world, i.e. that the agent’s actions can be sources of indeterminacy, then there will not be an underlying state space. On this approach, probability is not a purely epistemic matter reflecting uncertainty but is metaphysical. This ensures that there is no fact of the matter as to what would have happened instead. The fact that, say, N_1 (= the coin comes up heads) obtains at t' does not fix the counterfactual as to what would have happened had the agent performed the other action at t . In particular, it is then not possible to treat N_1 as a consequence of L_1 together with the underlying state, namely S_1 or S_2 , whereby this underlying

⁶As they point out, this is the analogue of applying the distinction between optimising and maximising to cases of risk, where an optimising choice set $C_{opt}(X) = \{x : \forall y(x \geq y)\}$ presupposes completeness whilst a maximising choice set $C_{max}(X) = \{x : \neg \exists y(y > x)\}$ is compatible with incompleteness (cf. Sen: 1997).

state determines the relevant counterfactual as to what would have happened had the agent opted for L'_2 (namely N_3 , i.e. the die would have landed 1-3, in the case of S_1 , whereas it would have been N_4 , i.e. die landing 4-6, in the case of S_2). Instead, the fact that N_1 rather than N_2 obtains will be a matter of chance. As a result, we will only have a probability assignment to the possible counterfactual states N_3 and N_4 . This means that there are only facts about what could have come about (with a certain probability) but not facts about what would have come about, which are required by state-by-state reasoning.

2.2 Expected-value fetishism?

Schoenfield has put forward a constraint, called LINK, that rules out the claim that L_2 is better than L_1 . (Bales, Cohen and Handfield have advocated a very similar constraint: COMPETITIVENESS (Bales et al.: 2014, p. 460). Hare identifies, though ultimately does not endorse, a closely related principle: RECOGNITION (Hare: 2010, p. 244).)

LINK: In cases in which considerations of value are the only ones that are relevant, if you are rationally certain that one option, A, will bring about greater value than the alternative option, B, you're required to choose A. If you are rationally certain that neither of the two options will bring about greater value than the other, it's not required that you choose A, and it's not required that you choose B. (Schoenfield: 2014, p. 267)

The thought is that since the outcomes under L_2 are non-comparable with those that would have resulted under L_1 , it is certain that no matter how the world turns out L_2 will not bring about greater value than L_1 , making it the case that the former lottery is not better than the latter and that both are permissible options.

This constraint is supported on the grounds that “[i]f expected value theory required us to make choices that we are certain would lead to no improvement in value, then expected value theory is imposing requirements that transcend what we actually care about: the achievement of *value*” (Schoenfield: 2014, p. 268). Theories that violate this constraint are criticised as being based on an ‘expected-value fetish’.

Contra Schoenfield, we are still concerned with value rather than expected value when judging that L_2 is better than L_1 . We are comparing the values of outcomes and are concerned with improvements in value. The difference between the state-by-state approach and the stochastic approach is not a difference in terms of the former being concerned with value whereas the latter is concerned with expected value. Instead, they differ in terms of how improvements are to be understood. This is because the accounts disagree over what the relevant alternatives for comparison are, in particular whether outcomes are to be evaluated only

relative to the obtaining state of nature or also relative to alternative equiprobable states of nature. The stochastic approach is not only concerned with how good things are relative to what would have come about but also relative to what could just as well have come about. As such, it is concerned with improvements in value and there is no commitment to an ‘expected-value fetish’.

2.3 Dominance and equiprobability

Stochastic dominance treats the comparison between L_1 and L_2 in the same way as that between L_1 and L_2'' , given that they involve the very same prospects, i.e. the outcomes and probabilities are the same and they only differ in terms of the assignment of outcomes to states of nature.

	N ₁ (50%)	N ₂ (50%)
L ₁	A	B
L ₂ ''	A ⁺	B ⁺

Bales, Cohen, and Handfield try to disassimilate the comparison between L_1 and L_2'' from that between L_1 and L_2 , arguing that only the former but not the latter operates by means of dominance reasoning. They argue that in the former case “the intuition is insensitive to the credences we have in the different possibilities”, but that in the latter case we can establish the ordering $L_2 > L_1$ “*only where our credences are perfectly balanced* between the relevant alternatives” (Bales et al.: 2014, p. 468).

This argument, however, is misguided. Dominance reasoning is operative in both cases. In the one case we are using state-by-state dominance, whereas in the other case we are using stochastic dominance which allows for permutations of equiprobable states. Even in the latter case, it does not matter what the credences/probabilities are *as long as the relevant states are equiprobable*. Relative rather than absolute probabilities matter and in this respect the dominance principles are exactly analogous.

State-by-state dominance reasoning likewise presupposes that the probability of the relevant state is preserved across the different possible actions. In this sense our credences also need to be perfectly balanced. Whereas in the case of state-by-state dominance we compare $\phi(N_i)$ with $\psi(N_i)$ and here have to presuppose that $P(N_i|\phi) = P(N_i|\psi)$, in the case of stochastic dominance we compare $\phi(N_i)$ with $\psi(N_j)$ and here have to presuppose that $P(N_i|\phi) = P(N_j|\psi)$. In each case, credences/probabilities of the alternatives being compared need to be perfectly balanced. The only difference is whether this comparison takes place within states or also across states.

The satisfaction of this equiprobability condition in the case of state-by-state comparisons is a non-trivial matter. In fact, the requirement that states have to be probabilistically independent from actions is in place precisely in order to

ensure that the equiprobability condition is satisfied. As Nozick showed, state-by-state dominance reasoning breaks down if the probability of a state depends on which action is performed (cf. Nozick: 1969). If it is more likely that N_i will eventuate rather than N_j if the agent ϕ 's rather than ψ 's, then ψ -ing can have greater expected utility than ϕ -ing even when ϕ -ing state-by-state dominates ψ -ing, i.e. even when $\phi(N_i) > \psi(N_i)$ and $\phi(N_j) > \psi(N_j)$, as long as the outcome that results from ψ -ing in the state that is made more likely by ψ -ing is better than the outcome that results from ϕ -ing in the state that is made more likely by ϕ -ing, i.e. $\psi(N_j) > \phi(N_i)$.

This means that outcomes under the same state will not be equiprobable when probabilistic independence is not satisfied and hence cannot be compared by means of dominance reasoning.⁷ Accordingly, we can see that both stochastic and state-by-state dominance are subject to the very same equiprobability condition, namely that outcomes can be compared on the basis of dominance reasoning only if the states in which they occur are equiprobable.⁸

3 Conclusion

We have seen that insisting on state-by-state reasoning is unduly restrictive. As long as the identity of the state of nature is evaluatively irrelevant, equiprobable states can be treated interchangeably. Accordingly, one should use stochastic dominance in comparing lotteries even when dealing with incomplete orderings that allow for non-comparable outcomes.⁹

⁷As Gibbard and Harper: 1978 have pointed out, the probabilistic independence condition can be interpreted in two ways, depending on whether one is using causal or evidential probabilities. Whilst the causal and evidential independence conditions can come apart in Newcomb-style cases, they coincide in the kinds of cases that are at issue in this paper.

⁸Once it is recognised that state-by-state reasoning is not appropriate when probabilistic independence is not satisfied, we can see that sameness of state of nature by itself is not of significance. The question then arises as to what makes it the case that such reasoning is appropriate when probabilistic independence is satisfied. Is it the fact that we are dealing with the same state of nature together with probabilistic independence (which, as we have just seen, amounts to an equiprobability condition), or is it rather equiprobability by itself that is doing all the work? The considerations about independent lotteries and about the evaluative irrelevance of the identity of states of nature adduced above suggest that it is equiprobability by itself.

⁹For helpful discussions, I would like to thank Bassel Tarbush. I am particularly grateful to two anonymous referees for their detailed and helpful comments, which have led to substantial improvements.

References

- [1] AUMANN, R. Utility theory without the completeness axiom. *Econometrica* 30, 3 (1962), 445–462.
- [2] BALCH, M., AND FISHBURN, P. Subjective expected utility for conditional primitives. In *Essays on Economic Behaviour under Uncertainty*, M. Balch, D. McFadden, and S. Wu, Eds. North-Holland Publishing Company, 1974, pp. 57–69.
- [3] BALES, A., COHEN, D., AND HANDFIELD, T. Decision theory for agents with incomplete preferences. *Australasian Journal of Philosophy* 92, 3 (2014), 453–470.
- [4] BROOME, J. Is incommensurability vagueness? In *Incommensurability, Incomparability, and Practical Reason*, R. Chang, Ed. Princeton University Press, 1997, pp. 67–89.
- [5] CHANG, R. The possibility of parity. *Ethics* 112, 4 (2002), 659–688.
- [6] FISHBURN, P. Subjective expected utility: a review of normative theories. *Theory and Decision* 13 (1981), 139–199.
- [7] GIBBARD, A., AND HARPER, W. Counterfactuals and two kinds of expected utility. In *Foundations and Applications of Decision Theory*, C. Hooker, J. Leach, and E. McClennen, Eds., vol. 1. D. Reidel, 1978, pp. 125–162.
- [8] HARE, C. Take the sugar. *Analysis* 70, 2 (2010), 237–247.
- [9] KRANTZ, D., LUCE, D., SUPPES, P., AND TVERSKY, A. *Foundations of Measurement*, vol. I – Additive and Polynomial Representations. Dover, 1971.
- [10] LUCE, D., AND KRANTZ, D. Conditional expected utility. *Econometrica* 39, 2 (1971), 253–271.
- [11] NOZICK, R. Newcomb’s problem and two principles of choice. In *Essays in Honor of Carl G. Hempel*, N. Rescher, Ed. D. Reidel, 1969, pp. 114–146.
- [12] SCHOENFIELD, M. Decision making in the face of parity. *Philosophical Perspectives* 28 (2014), 263–277.
- [13] SEN, A. Maximization and the act of choice. *Econometrica* 65, 4 (1997), 745–779.
- [14] VON NEUMANN, J., AND MORGENSTERN, O. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.