

Arif Ahmed

Book review: Lara Buchak // risk and rationality

**Article (Published version)
(Non-refereed)**

Original citation:

Ahmed, Arif (2016) *Book review: Lara Buchak // risk and rationality*. *British Journal for the Philosophy of Science - Review of Books* . ISSN 0007-0882


© 2016 The Author

This version available at: <http://eprints.lse.ac.uk/73693/>

Available in LSE Research Online: April 2017

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

Lara Buchak // Risk and Rationality

 bjpsbooks.wordpress.com/2016/08/22/lara-buchak-risk-and-rationality/

Reviewed by Arif Ahmed

Risk and Rationality

Lara Buchak

Oxford: Oxford University Press, 2014, £45 (hardback)

ISBN 9780199672165

Expected Utility

Subjective decision theory has two branches: descriptive and normative. The descriptive one says in general terms how actual people behave. The normative one says how rational ones do. Philosophers of decision theory focus mainly on the normative branch, economists and psychologists on the descriptive. But the basic classical theory applies the same abstract structure to both; and I'll start with that.

Very briefly then, Savage ([1972]) in effect distinguished: (i) a set \mathbf{S} of relevant possible states of the world; (ii) a set \mathbf{O} of possible outcomes of these states; (iii) a set $\mathbf{A} =_{\text{def.}} \mathbf{O}^{\mathbf{S}}$ of 'acts', this being the set of all functions from (i) to (ii); (iv) a binary relation \succ of 'strict preference' over the elements of (iii), together with a defined relation \succcurlyeq ('weak preference'), where $A_1 \succcurlyeq A_2 =_{\text{def.}} \neg(A_2 \succ A_1)$.

For instance and with some simplification, you are wondering whether to take your umbrella before going out into a world where it either is or is not raining and where all that matters to you is whether (D) you are dry and whether (U) you have to carry around an umbrella. Then, (i) the set of states is $\mathbf{S} = \{R, \neg R\}$, where R says that it is raining. (ii) The set of outcomes is $\mathbf{O} = \{DU, D\neg U, \neg D\neg U\}$. (iii) The set \mathbf{A} of acts is the set of (all) functions from \mathbf{S} to \mathbf{O} ; in this case \mathbf{A} has nine elements amongst which are the acts A_1 and A_2 , where $A_1(R) = A_1(\neg R) = DU$, $A_2(R) = \neg D\neg U$ and $A_2(\neg R) = D\neg U$. A_1 and A_2 are the actions of taking your umbrella and leaving it, respectively. (iv) Your strict preference over \mathbf{A} is a partial ordering of these acts that is typically (and in Savage's original story, analytically) revealed in your actual choices from amongst them. You have (for example) the preference $A_1 \succ A_2$ just in case you'd take your umbrella if leaving it at home were the only alternative.

Savage showed that \succ obeys certain axioms, nowadays called the Savage axioms, if and only if there is an associated subjective probability function, $Pr: \mathbf{S} \rightarrow [0, 1]$, and an associated utility function, $U: \mathbf{O} \rightarrow \mathbb{R}$, such that for any $A_1, A_2 \in \mathbf{A}$:

$$(1) \quad A_1 \succ A_2 \text{ if and only if } \sum_{s \in \mathbf{S}} Pr(s)U(A_1(s)) > \sum_{s \in \mathbf{S}} Pr(s)U(A_2(s)).$$

Moreover, Pr is unique and U is unique up to positive affine transformation. So any set of Savage-compliant preferences over the entire space $\mathbf{O}^{\mathbf{S}}$ of acts settles Pr altogether; and although it does not fix a zero for utility, it does completely determine ratios of utility intervals, that is, $(U(O_1) - U(O_2)) / (U(O_3) - U(O_4))$ whenever you are not indifferent between O_3 and O_4 .

Interpreting subjective probability as belief and subjective utility as desire, (1) makes a tight and interesting connection between one's belief and desire on the one hand and one's behaviour on the other. (1) says that the person is maximizing expected utility (EU): when choosing what to do in conditions of uncertainty about the state of the world, she is always selecting that act out of those available that according to her beliefs does best on average at satisfying her desires. And Savage's result tells us that this is happening if and only if the agent conforms to the Savage axioms, these being in effect as spelt out in the endpapers of (Savage [1972]).

Return now to the descriptive/normative distinction. Savage's axioms can be seen as descriptions of us—that is, of choices that actual people make, or as recommendations to us—that is, descriptions of choices that rational people make. For instance, one of the axioms implies that \succ is transitive: if you choose (to consume, for example) apples when bananas are the only alternative and bananas when cherries are the only alternative, then

you choose apples when cherries are the only alternative. As a description of actual choice, transitivity of strict preference is plausibly false.^[1] But a good deal may be said for it as a demand of rationality; at any rate a good deal more has been said for it as such.

Risk-Weighted Expected Utility

Risk and Rationality argues against Savage from a normative standpoint. In particular, although Buchak agrees that expected utility maximization in the sense of (1) is sufficient for rational preference over acts, she denies that it is necessary for this. She objects to what she calls the global neutrality of the Savage approach. Global neutrality means that improving the outcome of an act in some possible state of the world by some fixed increment of utility makes a difference to the overall value of the act that is independent of the outcome that the act originally yields in that state or in any other state.

For instance, suppose that the possible states of the world are S_1 , S_2 , and S_3 and that the following four gambles $A_1 - A_4$ on these states give you final levels of wealth (in millions of dollars) as specified in the following table.

| | S_1 | S_2 | S_3 |
|-------|-------|-------|-------|
| A_1 | 0 | 5 | 0 |
| A_2 | 1 | 1 | 0 |
| A_3 | 0 | 5 | 1 |
| A_4 | 1 | 1 | 1 |

Table 1. *Allais paradox*

Global neutrality implies that for any fixed S_1, S_2, S_3 , $A_1 > A_2$ if and only if $A_3 > A_4$. After all, the difference between A_1 and A_3 is only that one's utility in the event S_3 is increased by \$1M, and the difference between A_2 and A_4 is the same. So (1), in which global neutrality is implicit, itself implies that you prefer A_1 to A_2 if and only if you prefer A_3 to A_4 .

Note that the pattern of preferences $A_1 > A_2$ if and only if $A_3 > A_4$ is descriptively false for some choices of S_1, S_2 , and S_3 . For instance, if S_1 is the state that a randomly chosen integer in $[1, 100]$ is 1, S_2 the state that it lies in $[2, 11]$ and S_3 the state that it lies in $[12, 100]$, then most people do in fact prefer A_1 to A_2 but also A_4 to A_3 . This is the famous Allais paradox.^[2] From a normative perspective that mismatch between Savage's theory and actual practice may be neither here nor there.

Buchak proposes an alternative that she calls risk-weighted expected utility theory (REU). REU adds to Pr and U a third parameter, the risk function R , and characterizes the rational agent as one whose preferences satisfy the following condition. Suppose that of all the outcomes that are possible given an arbitrary act A and possible states S_1, \dots, S_n , the outcome $A(S_1)$ is the worst, $A(S_2)$ is the second worst, and so on. (We can always arrange this with an appropriate numbering convention.) Then the agent is rational if and only if there is a probability function Pr , a utility function U and a risk function R , this being an increasing function from $[0, 1]$ onto a subset of itself that includes both 0 and 1, such that for any acts A, A' :

$$(2) \text{ REU}(A) =_{\text{def.}} U(A(S_1)) + R(\sum_{2 \leq i \leq n} Pr(S_i))(U(A(S_2)) - U(A(S_1)))$$

$$+ R(\sum_{3 \leq i \leq n} Pr(S_i))(U(A(S_3)) - U(A(S_2))) + \dots$$

$$+ R(Pr(S_n))(U(A(S_n)) - U(A(S_{n-1}))).$$

$$(3) \text{ } A > A' \text{ if and only if } \text{REU}(A) > \text{REU}(A').$$

(See p. 53.)

Informally, REU has the following meaning. The agent evaluates a gamble by starting with its worst possible outcome: he will certainly do at least as well as that, if he takes the gamble. Then he adds to this quantity the least possible amount by which the gamble could make him better off than that, weighted not by the probability that he does at least as well as this, but by the R -weighted probability that he does. Then he adds to this quantity the least possible amount that the gamble could make him better off than that, again weighted by the R -weighted probability that he gets at least the third-worst outcome of the gamble; and so on.

If R is the identity function then REU-maximization coincides with EU-maximization. But otherwise the REU-maximizer is not an EU-maximizer. In particular, if R is convex to the x -axis (for example, $R(x) = x^2$) then the REU-maximizer exhibits risk-averse behaviour, and if R is concave to the x -axis (for example, $R(x) = \sqrt{x}$) then the REU-maximizer exhibits risk-seeking behaviour, even if her utility function for money is linear and so incompatible with risk-aversion or risk-seeking as classically understood. The reason is that if R is convex then the R -weights of very high probabilities, which correspond in (2) to the worse outcomes of the gamble, figure more largely in the overall REU of the gamble than do the R -weights of lower probabilities; and contrariwise if R is concave.

To see how this works in practice, consider a gamble A that offers a net gain of \$1 if this fair coin lands heads on its next toss and a net loss of \$1 if it lands tails; let A' be the 'gamble' that yields a sure gain/loss of \$0 in either case. Assuming linear utility for money, we clearly have $EU(A) = EU(A')$: the EU-maximizer would just as soon take the gamble as forego it. But let us now calculate $REU(A)$ on the supposition that $R(x) = x^2$. In accordance with (2) we let S_1 be the event that the coin lands tails and let S_2 be the event that the coin lands heads. So we have:

$$(4) \quad REU(A) = -1 + 2(0.5)^2 = -0.5.$$

$$(5) \quad REU(A') = 0.$$

So this REU-maximizer would prefer not to take the gamble and indeed would pay up to 50¢ to forego it, in spite of the fact—with which this preference is inconsistent in Savage's framework—that every additional cent matters to her at least as much as the last.

This example also illustrates the global sensitivity of REU-maximization: how much difference adding a little money to one of the outcomes of A makes, itself depends on whether the money is added to the bad or to the good outcome. We could persuade an REU-maximizer with $R(x) = x^2$ to take the gamble by adding 67¢ to its downside. But we should have to add \$2 to its upside to get the same effect. Again this reflects the fact that for convex R , the R -weighted probability of a likely event (in this case, doing at least as well as the bad outcome) more than proportionally dominates that of an unlikely one (that is, doing at least as well as the good outcome). And again it reflects preferences that EU theory finds irrational.

Buchak's argument falls into three parts:

(a) Chapters 1–2: Intertwined with the informal exposition of the theory itself is an informal case that an attitude towards risk exists, whether or not it is rational. The Allais paradox in [Table 1](#) illustrates this point: the preferences $A_1 > A_2$ and $A_4 > A_3$ over the acts there defined are inconsistent with the assumption that $R(x) = x$ for all x , as required by EU, but consistent with a suitable choice of convex R (p. 71).

(b) Chapter 3/Appendix A: The formal exposition of the theory proves a representation theorem that establishes about (3) an analogue of what Savage established about (1) that is, that it holds for some U , Pr , and R just in case the agent's preferences over acts satisfy a set of axioms that make no mention of U , Pr , or R ; moreover Pr and R are unique and U is unique up to affine transformation. The issue between EU and REU is then the issue over which set of axioms is rationally compelling. Chapter 3 lays out the REU axioms and contrasts them with the EU axioms (p. 100). Appendix A gives formal details of the proof itself.

(c) Chapters 4–7: Buchak defends REU-maximization against the following arguments: (Chapter 4) that the behaviour it seeks to capture can be interpreted as EU-maximization after all, given a suitable redescription of the states involved; (Chapter 5) that it endorses genuine deviations from EU-maximization that make the agent's preferences inconsistent; (Chapter 6) that it endorses preferences that lead the agent into sure loss, or into sure relative loss, in suitably arranged cases of sequential choice; and (Chapter 7) that it exposes the agent to bets that even a foresighted REU-maximizer would take separately but would reject if offered as a single package.

There is plenty to say about all of these arguments, which Buchak generally presents with care (in the sense of being scrupulous over details) and with candour (in the sense of admitting the genuine costs of REU theory that they sometimes identify). Here I'll focus on just two points, concerning dominance (Chapter 5) and diachronic choice (Chapter 6).

Dominance

Returning to the 'Allais' situation in Table 1, let us suppose that the agent prefers A_1 to A_2 and A_4 to A_3 , this being consistent with REU-maximization but not with EU-maximization. And for any act A and set of states E , let us call A -in- E the gamble which one takes oneself to be facing if one chooses A and then learns that some state in E is actual. (Such gambles are not defined in the Savage framework, but what I am saying here can be put in his terms at some expense of simplicity; see p. 160, Footnote 17.)

Now if we write $E =_{\text{def.}} \{S_1, S_2\}$ and $E^* =_{\text{def.}} \{S_3\}$, then we can ask now, before anyone's learning which of E and E^* is actual, what preferences the agent has amongst the A_i -in- E and A_j -in- E^* , for $i = 1, 2$. Nothing that I have said settles this altogether, but the following facts are clear and must be admitted on all sides:

(i) A_1 -in- $E > (\geq) A_2$ -in- E if and only if A_3 -in- $E > (\geq) A_4$ -in- E , since A_1 -in- E is the same gamble as A_3 -in- E , and A_2 -in- E is the same gamble as A_4 -in- E .

(ii) She is indifferent between A_1 -in- E^* and A_2 -in- E^* since these are the same gamble, namely, a sure thing of \$0.

(iii) She is indifferent between A_3 -in- E^* and A_4 -in- E^* since these are also the same, namely, a sure thing of \$1M.

Given these facts, it looks as though the preferences $A_1 > A_2$ and $A_4 > A_3$ jointly violate a very plausible principle of dominance. We know that either E or E^* is actual. Suppose that for some A, A^* , if E is actual then you prefer doing A . And if E^* is actual you prefer doing A . Then you ought to prefer doing A . More generally and precisely, the principle is this:

Event-Wise Dominance (EWD): Let $\mathbf{E} = \{E_1, \dots, E_n\}$ be any partition of \mathbf{S} . Then for any $A, A' \in \mathbf{A}$, if A -in- $E_j \geq A'$ -in- E_j for every $E_j \in \mathbf{E}$ then $A \geq A'$. If in addition A -in- $E_j > A'$ -in- E_j for some $E_j \in \mathbf{E}$ then $A > A'$.^[3]

Now in the present case the principle implies that the Allais preferences $A_1 > A_2$ and $A_4 > A_3$ must be jointly irrational. For if A_1 -in- $E > A_2$ -in- E , then it follows by (i) that A_3 -in- $E > A_4$ -in- E , and it follows from (iii) that A_3 -in- $E^* \geq A_4$ -in- E^* ; hence by EWD $A_3 > A_4$, contradicting $A_4 > A_3$. Alternatively, if A_2 -in- $E \geq A_1$ -in- E , then since by (ii) A_2 -in- $E^* \geq A_1$ -in- E^* , it again follows from EWD that $A_2 \geq A_1$, contradicting $A_1 > A_2$. So we seem to have an argument against the Allais preferences and by extension against generalizing EU-maximization to REU-maximization.

This is a well-known and straightforward objection to any view that allows for the rationality of the Allais preferences; and as such it has a bearing, but no special bearing, on Buchak. But it is worth pursuing her own line of response to this difficulty (at pp. 162–9) since that is both novel and interesting.

Buchak begins by rejecting EWD on the grounds that even if there is some partition of \mathbf{S} into events E, E^* such that A -in- E is at least as good as A' -in- E and A -in- E^* is at least as good as A' -in- E^* , there may be some strictly finer partition with respect to which an analogous condition does not hold.^[4] This is indeed what happens in the Allais case: neither A_1 nor A_2 is at least as good as the other with respect to every element of the partition $\{\{S_1\}$,

$\{S_2\}, \{S_3\}$ }, which is strictly finer than $\{E, E^*\}$. Accordingly, she endorses only a weaker form of dominance, one that only prefers A to A' on the condition that no such refinement effects any such reversal. This is:

State-Wise Dominance (SWD): If $A\text{-in-}S \succcurlyeq A'\text{-in-}S$ for every $S \in \mathbf{S}$, then $A \succcurlyeq A'$; if in addition $A\text{-in-}E > A'\text{-in-}E$ for some big enough $E \subseteq \mathbf{S}$ then $A > A'$.

Here the technical term 'big enough' denotes events that are not null in Savage's sense, that is, in effect those that have positive probability (pp. 88–9). SWD is just the application of EWD to the maximally fine partition. Clearly in the Allais case it does not apply for the reason that I just stated. State-wise dominance is therefore consistent with REU-maximization; but the question is whether the rational agent stops there or goes all the way to event-wise dominance.

Here is an informal case where EWD seems intuitively to recommend a course of action that is both rational and beyond the reach of SWD. Let X and Y be particular lottery tickets and suppose that for whatever reason you prefer X to Y . Now I offer you a choice between (A) a bet that awards X if the next toss of this fair coin comes up heads and some arbitrary fee Z if it comes up tails, and (A') a bet that awards Y if heads and Z if tails. It looks completely obvious that you should prefer A to A' whatever your basis for preferring X -for-sure to Y -for-sure, so long as you do have the latter preference. EWD delivers this result; but SWD does not, neither by itself nor in conjunction with the other principles of REU, which allow that a rational agent might notwithstanding all of this prefer A' to A .

More generally, the best reasons for finding dominance attractive in the first place are reasons for going all the way to EWD. Consider an argument for applying EWD to the Allais case that seems merely to extend an argument that many have found plausible in the context of Newcomb's problem.^[5] Suppose that you weakly prefer $A_1\text{-in-}E$ to $A_2\text{-in-}E$. Then you must weakly prefer $A_3\text{-in-}E$ to $A_4\text{-in-}E$. Now imagine that you must choose between A_3 and A_4 , and further that an accomplice asks a fee of \$1 to tell you in advance which of E and E^* actually obtains. Since you prefer $A_3\text{-in-}E$ to $A_4\text{-in-}E$, you know that if you paid the fee and learnt E you would happily take A_3 . Since you are indifferent between $A_3\text{-in-}E^*$ and $A_4\text{-in-}E^*$, you know that if you paid the fee and learnt E^* you would again happily take A_3 . So why bother to pay? Both possible bits of information make it optimal to take A_3 , so you might as well do so in advance of paying the fee and even if the accomplice is absent, as in the actual case. So if $A_1\text{-in-}E \succcurlyeq A_2\text{-in-}E$ then $A_3 \succcurlyeq A_4$, contrary to the stated Allais preferences. A similar argument shows that if $A_2\text{-in-}E > A_1\text{-in-}E$ then $A_2 \succcurlyeq A_1$, again contrary to those preferences.

From the perspective of REU-maximization the trouble with this argument is the step from (i) 'You would be rational to take A_3 over A_4 if you learnt that E ' and (ii) 'You would be rational to do the same if you learnt that E^* ' to (iii) 'It is rational to take A_3 *ex ante*'. *Ex ante* the REU-maximizer might consider that the risks associated with A_3 make it not worthwhile, even though this is not the case in any possible *ex post* scenario (that is, whatever she learns about E). For the EU-maximizer that cannot be true: if $A_3\text{-in-}E$ is at least as good as \$1M for sure then so is A_3 itself, since we can get to it from $A_3\text{-in-}E$ by shifting probability from S_1 and S_2 to a state that itself guarantees \$1M. But for the REU-maximizer things are different: even if $A_3\text{-in-}E$ is REU-preferred to \$1M for sure, A_3 itself might not be, because by comparison with $A_3\text{-in-}E$ it may not shift enough R -weighted probability from the worst outcome (\$0) to the second-best outcome (\$1M) to compensate for the shift in R -weighted probability from the best outcome (\$5M) to the second-best. In fact this will be the case if $Pr(S_2|E)$ is sufficiently high, $Pr(S_2)$ sufficiently low and R sufficiently convex. Still, that does not so much illuminate what is wrong with the dominance argument as locate the disagreement over it.

To pursue matters further, let us adjust the payoffs in [Table 1](#) so that the REU-maximizer with sufficiently convex R has, but for dominance reasons the EU-maximizer still cannot have, the strict preferences $A_4 > A_3$ before learning whether E and $A_3 > A_4$ after learning whether E (and whatever she learns). For instance, we might increase the payoff to $A_3\text{-in-}\{S_3\}$ by \$1. Then REU still regards it as permissible strictly to prefer A_4 to A_3 before learning whether E and to prefer A_3 to A_4 after learning this, and whatever it is that you learn.

For Buchak, the fact that you would take A_3 over A_4 if your information were enriched in either way (that is, on

learning E or on learning E^*) is still consistent with denying that you should take A_3 *ex ante*. This is because from your perspective *ex ante*, learning whether E risks putting you in a worse position: it risks putting you further from the truth over something that matters. In particular, if the true state of the world is S_1 then learning E increases your credence in the false proposition that it is S_2 ; this shift in credence is what makes the REU-maximizer choose A_3 upon learning E .

So while the agent [*ex ante*] thinks that there is nothing wrong with her future self after learning whether E —she is not being irrational, and she has not changed preferences—the agent does think that there is a chance of something being wrong with the information: namely, [not that it is false but] that it will (rationally) lead her to do something that is in fact worse for her. And, in particular, [*ex ante*] this information is not worth the risk. (p. 193)

But this puts the Allais-type REU-maximizer in a very strange position. Prior to learning whether E , she must think that that information is not worth acting on, but after learning whether E —and nothing else, and whether she learns E or E^* —she must think that the information is worth acting on. Buchak's view is that there is no incoherence here, because there is just no 'single perspective to which the agent has access, relative to which the information is good or bad'. She continues:

The agent at different times is simply in different positions from which to evaluate the instrumental value of an action. Therefore, my defence of the agent neither adopting her [*ex post* preference for A_3] at the earlier time, nor adopting her [*ex ante* preference for A_4] at the later time is that the preferences she has at each time are the appropriate ones at that time, and while she does not think that there is anything amiss in her perspective at the other time, she also does not think it better than her current perspective. But this is not inconsistent, because there is no fact of the matter about which perspective is better, apart from what is actually the case. (p. 196)

But the perspectives are at least informally distinguishable in a way that makes each of them seem worse from the other one. Before learning the truth-value of E the agent must prefer not to learn it prior to acting because she knows *ex ante* that learning it will lead her ('rationally') to realize the *ex ante* dispreferred act A_3 . But after learning the truth-value of E (and whatever this is) she must be glad that she learnt it prior to acting, or at least indifferent to whether she did, because she thinks then that before she learnt it she was (again 'rationally', according to REU) inclined to realize the *ex ante* dispreferred act A_4 .

That shift seems entirely irrational. The only thing that she learns is whether E itself is true or false; and how can such a bit of information make her change her mind on whether it is itself worth acting on? Prior to learning whether E the agent thinks that learning whether E is a bad means to securing her overall end (the maximization of money); but she also thinks then that (i) if she learns that E or (ii) if she learns that $\neg E$, she will think that it was a good means to that end. So she must herself think *ex ante* that in at least one of those *ex post* cases she is making some sort of mistake, and moreover that it is precisely in virtue of her departure from EU-maximization that this regrettable situation has arisen.

This brings me to my next point, which concerns what is most obviously regrettable about it. Again, the point may have force against any view that rationalizes the Allais preferences. But again, Buchak's response to it is both distinctive and of intrinsic interest.

Sequential Choice

When there is a foreseeable reversal of preferences over acts arising from the convexity of the risk function, it is possible to arrange cases of sequential choice in which the REU agent certainly does worse by their common lights than any EU-maximizer with the same associated Pr and U . In the Allais case, this is possible if we give

the agent the option to pay not to learn in advance whether E . More precisely, consider the following sequential choice problem.




Figure 1. Paying to eliminate a choice (see p. 188)

In this diagram time runs from left to right. Circles correspond to chance nodes—that is (in this case), revelations of facts over which the agent has no control. Squares correspond to choice nodes—that is, points at which she must act. Each path corresponds to some possible sequence of revelations and acts. The figure at the end of each path through the tree states the final payoff to that sequence of acts and chance events.

At the outset the agent, who knows that she will face a choice between A_3 and A_4 , is given the option of learning whether E in advance. She can choose to learn whether E , or to pay \$1 to avoid learning this in advance. In either case she then gets the choice between A_3 and A_4 , at which point she wins a prize as in rows 3 and 4 of [Table 1](#).

Consider an EU-maximizing agent Alice for whom $EU(A_4) > EU(A_3)$ at the outset. Her being an EU-maximizer means that $A_4\text{-in-}E > A_3\text{-in-}E$; and of course she is indifferent between $A_3\text{-in-}E^*$ and $A_4\text{-in-}E^*$. Thinking through what she will do in the future, she realizes that learning whether E can only affect her choice between A_3 and A_4 if what she learns is E^* ; but even then her choice has no effect on her payoff. So in either case it makes no difference whether she learns that E . So at the outset she declines to pay the \$1 and ends up with \$1M for sure.

Contrast Alice with an REU-maximizing agent Bob who has, as REU-maximization allows, $A_4 > A_3$ and $A_3\text{-in-}E > A_4\text{-in-}E$. Bob knows in advance that if he learns that E he will choose A_3 . Since from his initial perspective this is suboptimal, he pays \$1 not to learn whether E . At the second choice point he then takes A_4 and so ends up with $\$(1M - 1)$. So Alice is certainly doing \$1 better than Bob in virtue of her being, and his not being, an EU-maximizer; and this is something that everyone could have known in advance. This well-known and longstanding argument seems to represent a strong case for EU-maximization. Let me turn now to Buchak's distinctive

response.

Note first that the argument presupposes a 'sophisticated' approach to sequential choice, according to which the rational agent's present choice is informed by his expectation of his own actual or counterfactual future choices. In this case, it is Bob's expectation that if he learns E then he will choose A_3 that motivates him to pay not to learn whether E before choosing between A_3 and A_4 . There are alternatives to the sophisticated approach, the most normatively plausible being the 'resolute' approach on which the agent selects a plan and then somehow sticks to it despite a subsequent preference to deviate (McClennen [1990]). But it has seemed to many that the sophisticated approach is better, if only because the resolute approach is not generally feasible (Broome [1992]). In any case, Buchak's position is that REU is defensible even on the assumption of sophistication.

She says in this connection that although Bob certainly is doing worse than Alice in this problem, he is not doing any worse than he could have done, because he does not even have the option that Alice in fact takes, namely, A_4 after learning whether E (p. 189). Clearly, the fact that Bob is foreseeably doing worse than Alice should raise no doubts about Bob's rationality if we grant that Bob could not be doing as well as her.

But just as God's foreknowledge of Bob's actual choice does not make it any less a choice, so too Bob's own foresight of his hypothetical choice, of A_3 over A_4 if he learns that E , is no reason for him or anyone else to deny that it is (a) a choice of A_3 over (b) an available alternative A_4 that he is foreseeing. And whilst it is true that if and when the time comes to make this decision, he will in fact and foreseeably choose A_3 because he will in fact and foreseeably prefer it, it would only seem to follow from this that A_4 -after-learning-whether- E was all along unavailable on some incompatibilist sense of availability. The alternative, compatibilist sense of availability is roughly that an act is available to you if you would have done it had you wanted to do it. To my mind, this is the more relevant notion; but then Buchak's argument fails.

Buchak's response would be that sophisticated choice is committed to her favoured notion of availability. To see why, let us consider a simpler sequential decision problem in which there is no uncertainty about anything outside the agent's control (p. 175). Charlie has at t_1 the option of drinking a glass of wine: if he chooses to drink then he faces another option of drinking a second glass of wine at t_2 , but if he abstains at t_1 then he drinks no wine at all. At the outset Charlie prefers ultimately drinking one glass to abstention, and he prefers abstention to drinking two glasses. But if he drinks one glass then his preferences change: he prefers drinking two glasses to drinking one glass.

The sophisticated approach to this problem has Charlie bear in mind from the outset that his preferences over final outcomes will change if he takes one glass of wine at t_1 . If he does that then he ends up drinking two glasses. So sophistication requires that Charlie abstain at the outset, for a final outcome of no drinking.

Next, notice that given a one-off or 'synchronic' choice between the three options of abstention, one glass and two glasses, Charlie will make a different choice despite having the same *ex ante* preferences over these outcomes. In this case a sophisticated approach to sequential choice will (like any other approach to sequential choice, which must agree with it over this degenerate 'sequence') recommend that Charlie take just one glass.

So sophisticated choice violates a principle due to Hammond that Buchak labels:

Only Logical Consequences (OLC): If two decision problems have the same logically possible consequences, and the plan the agent chooses in one results in consequence c , then the plan he chooses in the other problem must also result in consequence c . (p. 178)^[6]

The appeal of OLC is that it is supposed to follow from (Hammond's explication of) consequentialism: if the agent chooses different outcomes in problems with the same logically possible outcomes then he cares about something other than the outcomes, that is, something other than the consequences of his actions. And sophisticated choice violates this principle because both the sequential and the synchronic versions of Charlie's problem have the same logically possible consequences, but (as we just saw) sophistication has different consequences in the two cases.

According to Buchak, advocates of sophisticated choice can only occupy this position if they motivate the denial of OLC on the following grounds: that decision problems with the same logically possible consequences might still differ over which outcomes are really available (p. 180).

Returning to the sequential problem in [Figure 1](#), sophisticated choice can therefore endorse Bob's paying to avoid the information about E and then choosing A_4 . This policy certainly loses money relative to the logically possible alternative of taking the information and then taking A_4 , but the latter policy is simply not available to Bob at any time. Not *ex ante*, because he cannot then prevent his possible future self from choosing A_3 given the information whether E . And not *ex post*, because he cannot then prevent his actual past self from having chosen to forego the information.

But sophisticated choice does not proceed by 'eliminating the consequences that are not available': it proceeds by eliminating the consequence that will in fact not occur. To say that they will not occur just because of the agent's present or future preferences is to say that they are available but not in fact availed of. The policy of taking the information and then taking A_4 certainly is available to Bob (and also to Alice) *ex ante*.

But then how can an advocate of sophisticated choice demotivate OLC, with which he certainly does conflict in Charlie's case? The answer is that OLC is not in general a consequence of consequentialism: an agent may care only about the outcomes of any decision problem and yet violate OLC if his ranking of final outcomes changes across problems with the same logically possible outcomes. This is what happens to Charlie: in the initial, sequential problem there is a decision node at which he prefers ultimately drinking two glasses over ultimately drinking one, whereas in the synchronic version he prefers one glass to two at the only decision node that that problem contains. In both cases Charlie cares at all times only about the consequences, but might rationally end up with different outcomes in the two cases even in the absence of uncertainty. So sophistication may plausibly be combined with consequentialism without denying the availability, in a normatively relevant sense, of outcomes that foreseeably will not be realized.

From the sophisticated perspective therefore, Bob has as much reason *ex ante* to rue the endogenous preference change over acts wrought by REU-maximization as does Charlie to rue the exogenous preference change over outcomes that would arise from the first glass of wine. Either such change entails that some outcome is both genuinely available at the outset and actually preferred at all times to the one that is actually realized. In short, a sophisticated consequentialist should see REU-maximization as imposing the same costs on an agent as Charlie's weakness of will, namely, that it leads to gratuitous loss. This casts a normatively unfavourable light on REU-maximization if anything ever casts such a light on any decision rule.

Whatever its descriptive adequacy, this being a question that Buchak largely avoids and which I have sidestepped altogether, I therefore remain uncertain that REU-maximization has the normative advantages that this book claims for it.

Conclusion

I should finally make explicit three points about the book that probably are not even implicit in the above.

First, the level of argument is very high and arguments are pursued in detail over a wide range of sub-topics. Here I have followed two interesting threads of Buchak's argument; but I could have done the same for a dozen others.

Second, the book as a whole contains a wealth of useful material, of which Buchak's exposition is thorough and extensive without straying off the point. Almost anyone who cares about the subject of this book is likely to learn much from it that is worth learning. It is all in all an excellent addition to the literature.

Finally and despite my reservations about it, REU-maximization deserves further investigation. An especially fertile point may be the analogy between the problem of distributing concern over one's possible future selves and that of distributing goods over a population of actual people (see pp. 167–8). In that analogy, REU-maximization with a convex R prioritizes fairness above individual welfare, excessively so from the EU perspective. In itself this fact is no argument for anything; but pushing similar parallels as far as they can go may

cast interesting light on the differences between society and the individual as objects of rational concern.

Arif Ahmed
Department of Philosophy
University of Cambridge
ama24@cam.ac.uk

References

- Ahmed, A. [2014]: *Evidence, Decision and Causality*, Cambridge: Cambridge University Press.
- Allais, M. [1953]: 'Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école Américaine', *Econometrica*, **21**, pp. 503–46.
- Broome, J. [1992]: 'Review of McClennen 1990', *Ethics*, **102**, pp. 666–8.
- Hammond, P. [1988]: 'Consequentialist Foundations for Expected Utility', *Theory and Decision*, **25**, pp. 25–78.
- McClennen, E. [1990]: *Rationality and Dynamic Choice*, Cambridge: Cambridge University Press.
- May, K. O. [1954]: 'Intransitivity, Utility, and the Aggregation of Preference Patterns', *Econometrica*, **22**, pp. 1–13.
- Nozick, R. [1990]: 'Newcomb's Problem and Two Principles of Choice', in P. Moser (ed.), *Rationality in Action: Contemporary Approaches*, Cambridge: Cambridge University Press, pp. 207–34.
- Pollock, J. L. [2010]: 'A Resource-Bounded Agent Addresses the Newcomb Problem', *Synthese*, **176**, pp. 57–82.
- Savage, L. J. [1972]: *The Foundations of Statistics*, New York: Wiley.

Notes

[1] See, for example, (May [1954]).

[2] (Allais [1953]).

[3] (p. 164, Footnote 23). Fatalism would seem to generate counterexamples to EWD as here stated (and indeed to the weaker principle SWD that I am about to explain). But it is possible to get around this by placing suitable general restrictions on the available event-partitions in a way that raises no special difficulties in the present case (p. 91).

[4] Cf. (McClennen [1990], pp. 77-80).

[5] (Pollock [2010], pp. 70-3; cf. Nozick [1990], pp. 209-10). For my own reservations about the argument as applied to the Newcomb case, see my ([2014], Section 7.4.2).

[6] See (Hammond [1988], pp. 37f).