Title
Replies to Commentators

Journal
Philosophical Studies, 174(9)

ISSN
0031-8116

Author
Buchak, Lara

Publication Date
2017-09-01

DOI
10.1007/s11098-017-0907-4

Peer reviewed

Penultimate draft—the final publication is available at Springer via http://dx.doi.org/10.1007/s11098-017-0907-4

Replies to Commentators

Lara Buchak

ABSTRACT

I reply to two commentaries—one by Johanna Thoma and Jonathan Weisberg and one by James M. Joyce—about how risk-weighted expected utility theory handles the Allais preferences and Dutch books.

1. Introduction

I am honored to be able to respond to the two excellent critiques of REU theory in this volume, put forward by Johanna Thoma and Jonathan Weisberg and by James Joyce. Both essays criticize how REU theory handles the Allais preferences, but for different reasons. Thoma and Weisberg argue that the Allais preferences are difficult to characterize as REU-maximizing, once we notice that money itself is really a gamble over various possible outcomes. Joyce argues that the Allais preferences can be characterized as EU-maximizing, in combination with two (irrational) psychological tendencies: the endowment effect and loss aversion. In a different vein, Joyce criticizes my argument that REU-maximizers are not subject to Dutch books.

2.1. Thoma and Weisberg's critique: No plausible model of REU theory can handle the complex Allais problem

Thoma and Weisberg point out that money is not valuable for its own sake. Instead, a given amount of money is a gamble over various ways the world could turn out. For example, if I get \$200 and plan to spend it at a particular store, various items might be available or on sale, so \$200 is a gamble over the various possible items. Or, I may use \$200 to by a new bicycle, which has some chance of leading to increased fitness but some chance of leading to injury, so \$200 spent on a bicycle is a gamble over various health states. This point is even more stark with larger sums of money. If I receive \$1M, there will be a very large range of items newly available to me. If I use \$1M to buy a particular house, this could increase my well-being in various ways (I might develop a life-long friendship with the neighbor, my children might receive an excellent education in the school district), but it could also decrease it (there might turn out to be lead in my soil, I might become the victim of a break-in).

More generally, let's say we face a gamble with some known consequences. We can make a distinction between consequences that are outcomes—objects that are valuable for their own sake—and consequences that are further gambles over outcomes. Call the latter 'sub-gambles,' since they are gambles within a larger gamble. This distinction is particularly important in REU theory, since the risk-weighting applies to probabilities of outcomes and cannot directly be applied to probabilities of sub-gambles. This is because sub-gambles don't have fixed utility values, which is because the value a sub-gamble contributes to a gamble in which it is embedded depends on the gamble.[1] By contrast, in EU theory, a sub-gamble contributes the same value in all gambles in which it is embedded. Thus, unlike the EU-maximizer, the REU-maximizer who makes calculations using the values of sub-gambles will arrive at different (and erroneous) recommendations from the REU-maximizer who makes these calculations using the values of the outcomes themselves.

One way to formulate this point is in terms of the grand-world and small-world decision problems. The grand-world decision problem is the problem in which all relevant states are distinguished and all outcomes are fully specified, whereas the small-world decision problem is a simplification of this problem, where some states are lumped together even though they may lead to differently-valued outcomes. Normative decision theory has typically been taken to be about how one should choose in the grand-world decision problem, and both EU and REU adhere to this. But because EU calculations using sub-gambles are the same as those using outcomes, EU theory can more easily *also* be a theory about how to choose in the small-world decision problem, in the following sense. If we know the expected utility of a sub-gamble, then we do not need to know, before we perform EU calculations, whether the value of this sub-gamble consists in an outcome (or outcomes) of a single value, or consists in a distribution over outcomes with widely diverging value: the calculations come out the same either way.[2] Although this means the REU-maximizer needs to know more facts before making a decision, I take this to be a positive feature of REU theory, because it more faithfully reflects what is of ultimate relevance in decision-making.[3] However, Thoma and Weisberg argue that this feature may pose a problem for my interpretation of the Allais paradox.

The original statement of the Allais choices are in terms of monetary values:

---

[1] See Buchak (2013: 160-169).

[2] See Briggs (2015) for a helpful example.

[3] See Buchak (2015: 16-19).

L_1 (C): $5,000,000 with probability 0.1, $0 with probability 0.9.
L_2 (D): $1,000,000 with probability 0.11, $0 with probability 0.89.

L_3 (B): $1,000,000 with probability 0.89, $5,000,000 with probability 0.1, $0 with probability 0.01.
L_4 (A): $1,000,000 with probability 1.

I've included Thoma and Weisberg's labels (C, D, B, A) for each choice, but I will stick with the "L_i" terminology for the sake of agreement with *Risk and Rationality* and with Joyce's response.

Call the above-listed choices—where outcomes are taken to be amounts of money—the *simple* Allais choices. People typically prefer L_1 to L_2 and L_4 to L_3. These are the preferences that EU theory cannot reconstruct, since there is no utility assignment to $0, $1M, and $5M according to which L_1 has a higher EU than L_2 and L_4 has a higher EU than L_3. (The outcomes are technically \$$i$ + 0, \$$x$ + $1M, and \$$i$ + 5M, where \$$i$ represents initial wealth, but I will drop the initial wealth variable in what follows.) But REU theory *can* reconstruct them. An individual will have the typical preferences if and only if his risk-function is convex (risk-avoidant) and the ratio of [u($5M) – u($0)] to [u($1M) – u($0)] is between two particular values, where these values are determined by his risk function.[4]

Thoma and Weisberg point out, however, that it would be a mistake to assign utility directly to amounts of money, since each monetary amount is a sub-gamble, not an outcome. They suggest that we think of the three monetary amounts involved each as a normal distribution centered around 'the' utility value of the amount in question (we will write c(x) to represent the center of the distribution), each with some variance σ. Let us label the complex gambles that result from this suggestion—the gambles that might yield *sub-gambles* $0, $1M, and $5M—$L_1^*$, L_2^*, L_3^*, L_4^*. These are the *complex* Allais choices. Thoma and Weisberg's assumption is that the complex Allais choices correctly represent the decision problem that individuals face, and thus the task is to reconstruct the complex Allais preferences (for L_1^* over L_2^* and L_4^* over L_3^*).

Now, the complex Allais preferences still cannot be squared with EU-maximization (since there is no difference in EU theory between the simple and the complex calculations), and the goal of Thoma and

[4] Buchak (2013): 71.

Weisberg's criticism is not to save EU theory from the counterexample. What Thoma and Weisberg instead show that it is difficult to square the complex preferences with REU theory.

Here are some of their results. Assume that c($0) = 0, c($1M) = 1, and c($2M) = 2. Then if we assume that $\sigma \geq 0.2$, there is no value of x for which $r(p) = p^x$ recovers the complex Allais preferences. For $\sigma = 0.1$, the complex Allais preferences can be recovered for values of x between 2.05 and 2.9 (i.e. between $r(p) = p^{2.05}$ and $r(p) = p^{2.9}$), and for $0.1 < \sigma < 0.2$, the preferences can be recovered within an increasingly narrow range of x-values. However, Thoma and Weisberg argue that setting $\sigma < 0.2$ makes the variance implausibly small, because the possible values of the $0-gamble and the $1M-gamble won't overlap enough. For example, if $\sigma = 0.1$, then an individual would have to assign probability 0.9999999999999 to the claim that his life would be better with the $1M than without it, which is implausible.

Thoma and Weisberg think of the simple Allais choices as corresponding to the small-world decision problem, and the complex Allais choices as corresponding to the grand-world decision problem. They conclude that since the complex Allais choices cannot be reconstructed without implausible assumptions, REU theory loses some of its initial motivation as a suitable grand-world theory. They further conclude that REU theory cannot be re-cast as a small-world theory (a conclusion I agree with), and thus that REU theory is not a promising theory of choice in either vein.

This is an important and instructive challenge for REU theory. There are two ways for a proponent of REU theory to respond, and I turn to them now.

2.2. First response: We can resolve the problem by tweaking other variables

In *Risk and Rationality*, after deriving the conditions that reconstruct the simple Allais preferences, I gave an example of a utility function and a risk-function that meet these conditions. The risk-function was $r(p) = p^2$, and the utility function assigned u($0) = 0, u($1M) = 1, and u($2M) = 2. Unsurprisingly, this is the utility function that Thoma and Weisberg use for their investigation. However, if we are talking about actual preferences, then the utility function mentioned above is likely too sloped to represent the preferences of most actual people: a 2:1 ratio of [u($5M) – u($0)]:[u($1M) – u($0)] is too high. Compare the given utility function, for example, to a standard example of a diminishing marginal utility function of total wealth, u($*i* + $*a*) = ln($*i* + $*a*), where *i* represents one's initial wealth and *a* the amount of money added to one's total fortune by the present gamble. For initial wealth levels under about $340,000, the ratio between the value of adding $5M rather than $0 to adding $1M rather than $0 is less than 2. And it is easier to reconstruct the complex Allais preferences the lower the ratio between these two values. For example, although $r(p) = p^2$ and variance 0.1 does not reconstruct the preferences for the 2:1 utility ratio, it does for the 1.3:1 utility ratio.[5]

That being said, this doesn't touch Thoma and Weisberg's point that a variance of 0.1 is implausibly small because it means the individual is nearly certain that his life with the $1M is better than his life without it. And, as it turns out, a variance this small is still necessary to reconstruct the preferences for the 1.3:1 utility ratio.

Is there a way to reconstruct the complex Allais preferences without the implausible implication, while being faithful to how a real-world decision-maker should view $1M? One possibility that Thoma and Weisberg explore is to vary the variance associated with each of the monetary amounts. They note that "those with $1M in the bank are less vulnerable to many of life's setbacks: they are not as easily ruined as the rest of us,"([11]), and following this line they consider the possibility that $1M and $5M are gambles with progressively smaller variances, since more money provides more security. They note that the complex Allais preferences are recoverable, though only for certain highly specific combinations of variances and risk-function exponents, where the latter values make the individual implausibly risk-avoidant.

But this isn't quite the right way to think about the added security of $1M. It is not that there are fewer overall possibilities represented by receiving $1M or that probability is more concentrated around the mean—indeed, receiving $1M opens up a larger range of possible outcomes—but rather that receiving $1M makes the worst possibilities much less likely. It provides security in the sense of making the probability associated with lower utility values smaller and smaller. The utility of $1M is concentrated around a high mean with a long tail to the left: things likely will be great, though there is some small and diminishing chance they won't be very good.[6] Similarly, the utility of $0 is concentrated around a low

[5] See Appendix B.

[6] Another possibility that Thoma and Weisberg don't explore is the possibility that variance *increases* as the amounts of money increase. This seems fairly plausible, since receiving $1M or $5M opens up a large range of possibilities that weren't present before. (And, again, the security added by the higher monetary value is security in the sense of eliminating or lowering the probability of the bad possibilities, not in the sense of narrowing the overall range of possibilities.)

mean with a long tail to the right: things likely will be fine but not great, though there is some small and diminishing chance they will be increasingly great. In other words, \$1M (and \$5M) is a gamble with negative skew, and \$0 is a gamble with positive skew:

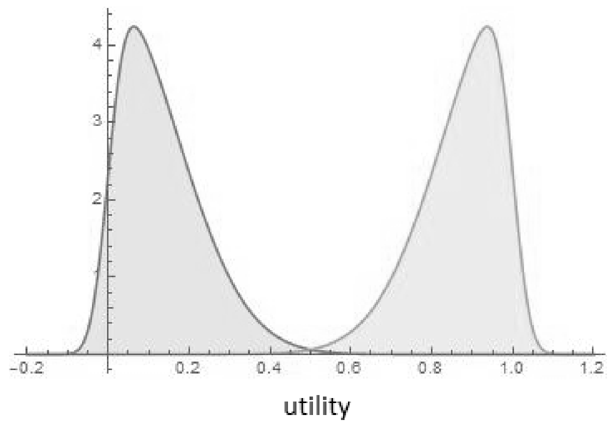

Figure 1: \$0 Gamble and \$1M Gamble

When we add in skewness, we can recover the Allais preferences for a wider range of variances, and without the implausible upshot that there is an astronomically low chance that \$0 is worse than \$1M. For example, if we want to recover the complex Allais preferences for a 1.3:1 utility ratio and $r(p) = p^2$, then we can set the variance of all the gambles up to 0.17,[7] the skew of \$0 at 5, and the skew of \$1M and \$5M at -5. (We can also recover the complex Allais preferences for these utility and risk function with less skew, at lower variance.) Furthermore, with these values, there is more overlap in the utility that \$0 and \$1M might deliver, and thus more of a chance (albeit, and correctly, still a small one) that the former gamble will be better than the latter: in this example, there is a 0.003 probability that the \$0 gamble will deliver more than 0.5 utils, and a 0.003 probability that the \$1M gamble will deliver less than 0.5 utils (Figure 1).[8]

These possibilities allow, contra Thoma and Weisberg, that the complex Allais preferences could be rational for a not-overly-risk-avoidant REU maximizer who makes plausible assumptions about the gambles involved. The lesson for the debate about REU theory is that if we want to consider the complex Allais choices, there are many complexities to consider, and we need to include them all.

[7] See Appendix C. This was roughly the maximum variance that reconstructed the preferences given a 1.3:1 utility ratio and skew 5/-5.

[8] See Appendix C.

However, it is worth noting that these possibilities (like Thoma and Weisberg's) are all fairly speculative, empirically speaking: we don't really know how most people who have the Allais preferences view the amounts of money involved. The second response side-steps the need to know this and focuses on what typically goes on in thought-experiments.

2.3. Second response: Thought-experiments define the grand-world problem

The second response to Thoma and Weisberg holds that individuals facing the Allais choices take them at face value: decision-makers really are facing the simple Allais choices.

Consider the usual role of thought-experiments. When we elicit intuitions in thought-experiments, we typically ask people to abstract away from details that are not explicitly mentioned, even if they are ordinarily features of cases like the ones in question. For example, when asking whether an individual will 'pull the switch' in a trolley problem, we count responses such as 'I will jump onto the tracks myself' or 'I will do whatever takes more time so that there's a chance someone will notice the oncoming trolley and move off of the tracks' illegitimate: if someone were to give this type of response, they haven't accepted the setup of the problem.

Similarly, we can hold that the request to consider your preferences in the Allais example is a request to abstract from potential complications—to not really take seriously the chance that \$1M might be worse than \$0—and to treat the monetary outcomes as stand-ins for some goods that are valuable for their own sake. Thus, we can hold that while the grand-world decision problem is what an individual should ultimately be solving, individuals given the Allais choices are thinking of them as *simple grand-world* choices.

Support for this point comes from the fact that typical reasoning for the Allais preferences concerns the fact that taking L_4 risks nothing, whereas all the other gambles come with a much greater risk.[9] Individuals facing these hypothetical choices appear not to think of \$1M as a gamble. Other support comes from studies in which the Allais preferences are replicated over non-monetary consequences, such as grades or health.[10] Consequences in these domains are more plausibly outcomes in the strict sense:

---

[9] See MacCrimmon (1968), Moskowitz (1974), and Slovic and Tversky (1974).

[10] See Moskowitz (1974) and Oliver (2003), respectively.

unlike money, they are not primarily means to particular ends (though perhaps the case of grades is somewhere between money and health on this scale).

Perhaps decision-makers facing the Allais choices *ought* not to see them as the simple choices. But as it turns out, this claim would not be a strike against REU theory. To see this, consider the particular role of the Allais preferences in the argument for REU theory. The Allais case is relatively unique, and it concerns a choice that (presumably) next to no one ever faces. The main point of the Allais case in the argument isn't that it is so important to capture *these* preferences—as if any theory that missed them would fail to capture the vast number of Allais choosers out in the wild—but rather that the intuitive nature of these preferences reveals that there is some aspect of decision-making that EU theory is missing.

To see this, let us examine the reasons that support the Allais preferences: in the choice between L_1 and L_2 the minimum is the same so one might as well go for the (nearly identical) chance of the higher amount, but in the choice between L_3 and L_4 the fact that the latter has a much higher minimum swamps all other considerations. These reasons are reflected in the formal structure of REU theory. If a decision maker uses reasoning that cites the importance of what happens in relatively worse states rather than relatively better states, then this has a formal correlate in REU (higher weight is given to states of relatively lower value), and the resulting preferences will be captured by REU-maximization with a risk-avoidant risk function.

The point is that a certain kind of reasoning rationalizes the simple Allais preferences and is captured by REU maximization. Thoma and Weisberg argue that REU-maximization cannot capture the complex Allais choices, where the complex choices are choices to which the reasoning just mentioned does not apply. So, if people 'really' face the simple choices, then their reasoning is correct and REU captures it. If people 'really' face the complex choices, then the reasoning in favor of their preferences is misapplied, and REU does not capture their preferences. Either way, the point still stands: REU maximization rationalizes and formally reconstructs a certain kind of intuitive reasoning, as seen through its ability to capture highly idealized gambles to which this reasoning is relevant.

So, Thoma and Weisberg's results do nothing to dismantle the conclusion that REU theory explains why a certain kind of intuitive reasoning can be seen as consistent reasoning towards the means to one's ends, and equivalently the conclusion that it better captures intuitive preferences among certain highly idealized gambles.

3. Joyce's first critique: REU-maximization is not needed to reconstruct the Allais preferences

James Joyce, however, explicitly critiques this conclusion. He claims that the Allais preferences (including the reasoning that produces them) can be explained by a package of tendencies of actual decision-makers that both Joyce and I agree are irrational. These are a tendency to over-value losses, in the sense of assigning higher negative utility to a loss of a given amount than positive utility to a gain of that amount, relative to a perceived status quo; and a tendency for the setup or description of a decision problem to influence what one counts as the status quo in particular ways. Call this package of tendencies *unstable outcome-preferences*.

I don't want to argue with Joyce's claim that these tendencies could give rise to the Allais preferences. Indeed, individuals who have the Allais preferences might have them for a variety of reasons. If we want to discover the reasons for a particular individual's Allais preferences, we will have to know more than the preferences themselves.

As it turns out, there is some empirical evidence about why people have the Allais preferences. The earliest empirical study on the Allais paradox is due to MacCrimmon (1968), who asks participants first to determine their preferences in the Allais choices, and then to choose between two prepared answers about why one set of preferences was better, according to which answer was more logical.[11] One answer corresponded to choosing L_2 and L_4, and the reasoning given was that consistent with EU-maximization (Savage's (1954) 'sure-thing principle'); the other corresponded to the Allais preferences and the reasoning given was that L_4 was safer than L_3 but that both L_1 and L_2 involved a high and nearly identical chance of the bad payoff, so one might as well go for a chance of the higher payoff (L_1). There were two experiments of this type: in the first, 40% of all subjects selected the Allais choices and 80% of all subject selected the Allais reasoning, and in the second, roughly the same percentage of all subjects selected the Allais choices and 50% of all subjects selected the Allais reasoning (the main difference between the two experiments was that in the second the presentation of alternatives was more qualitative).

The study that explicitly considers whether the Allais preferences come from rank-dependent probability weightings or unstable outcome-preferences is due to Oliver (2003). In a study involving the Allais choices over health outcomes rather than money, Oliver asked individuals to make choices between health analogues of L_1 and L_2 and also between health analogues of L_3 and L_4; he also asked individuals

[11] See also Moskowitz (1974), Slovic and Tversky (1974), and MacCrimmon and Larsson (1979).

to give their reasons for their choices. 14 of his 38 subjects displayed preferences analogous to the Allais preferences.[12] Of the 7 who gave reasons for both choices, 2 gave reasons that suggest rank-dependent probability weighting, 2 gave reasons that suggest unstable outcome-preferences,[13] and 3 gave reasons that suggest a hybrid of the two. Granted, the study was small and more work would need to be done, but the findings suggest that my preferred explanation and Joyce's are each at work in some individuals' preferences.

Again, *if* an individual's reason for holding the Allais preferences is that he cares more about how things go in relatively worse states, *then* his preferences are well-captured by REU-maximization. And it is inarguable that at least some people have the Allais preferences for these reasons. So, with both critiques, the point still stands: REU-maximization formalizes the kind of informal reasoning that is both plausible in its own right and could give rise to the (simple) Allais preferences.

4. The general picture: theory and data

Let me say something more general about the picture I am working with. What we are trying to come up with is a decision theory that explains the data. The data consists in: individuals' stated preferences in the lab; their intuitive preferences in thought-experiments; the reasoning behind preferences in both settings; the robustness of these preferences in the face of alternative reasoning; these preferences in combination with self-reports of individuals' utility functions; and individuals' actual behavior. To explain the data is to do two things: to explain the structure of rational preferences (the normative component) and to explain the ways in which people systematically err from adhering to this structure (the descriptive component).

My claim is that the normative component is given by REU-maximization, and the descriptive component includes the tendencies mentioned above, and whatever else empirical psychology tells us is a widespread phenomenon. There is empirical evidence besides the Allais paradox that seems to show that rank-

[12] 20 of the participants violated EU-maximization, which Oliver notes is consistent with studies that have used monetary outcomes.

[13] These reasons were also consistent with anticipated regret and anticipated disappointment; the main point is that they do not suggest probability weighting.

dependent probability weighting is descriptively robust.[14] Thus, there seem to be strong reason for including rank-dependence in either the normative or descriptive component.

We then face a separate question of whether to count rank-dependent probability weighting as normative or descriptive.[15] There are both empirical and philosophical reasons for counting it as normative. On the empirical side, we have evidence that decision-makers *treat* rank-dependence as normative: they stick with the Allais preferences (bolstered by rank-dependent reasoning) even after being given the reasoning that 'should' lead one to reject these preferences; this is in contrast to how decision-makers treat other deviations from EU-maximization, such as intransitivity.[16] Thus, if we hold that it is a mark in favor of a decision theory that it can rationalize what people take to be their reasons—and take to be their reasons even upon reflection—then we have a reason to include rank-dependence in the normative component.

What do the philosophical arguments say about whether rank-dependence, and REU-maximization in particular, is normative or descriptive? Since EU-maximization is a special case of REU-maximization, in order to make my case, I need to show that there is nothing superior about EU-maximization in particular, and that non-EU-maximizers don't fall prey to something we can all agree is irrational. Joyce's final critique is that my argument for this claim fails, and it is to this critique that I now turn.

5 Joyce's second critique: REU-maximizers do not have an adequate response to Dutch books

---

[14] See especially Starmer (2000), who surveys the alternatives to EU-maximization and concludes that the evidence in favor of both probability weighting and loss aversion is strong, so that rank-dependent theories (including those that build in loss-aversion) are the most descriptively promising. For other surveys of alternatives to EU-maximization (including rank-dependent theories) and discussions of how these fit with empirical results, see Machina (1987), Camerer (1989), Sugden (2004), and Schmidt (2004). A different kind of result comes from an experiment by Abdellaoui, Barrios, and Wakker (2007), who asked individuals to determine their utility functions on the basis of introspection and also asked individuals to determine their preferences between particular gambles; the introspected utility functions agreed with those derived from rank-dependent utility theories but not with those derived from expected utility theory.

[15] We also face the question of whether to count tendencies like those mentioned above as part of the normative or descriptive component. I argue that these should not be included in the normative component in Buchak (2013: 74-81).

[16] See MacCrimmon (1968), Moskowitz (1974), and Slovic and Tversky (1974). However, MacCrimmon finds that after discussion with an experimenter, a substantial portion of subjects endorse the EU-conforming preferences and reasoning. See MacCrimmon (1968) for evidence that people do want to conform their preferences to transitivity.

Joyce claims that the normative foundations of REU theory are unsupported because (contrary to what I've argued) REU theory does not have a plausible response to the Dutch book argument. The Dutch book argument is an attempt to show that individuals who do not price bets according to EU-maximization can be made to buy or sell combinations of bets in such a way that they incur sure loss or forgo sure gain.

To illustrate, consider Joyce's example of Jacob, who has $u(x) = \$x$ and a $r(p) = p^2$. He will visit two bookies. Where <a, b> denotes the bet that adds $\$a$ to Jacob's fortune if a particular event of probability 0.3 occurs and that adds $\$b$ to Jacob's fortune if it the event does not occur, the first bookie offers the bet <1, 0> for 10¢ and the second offers the bet <-1, 0> for -50¢ (i.e. he offers to buy the bet <1, 0> from Jacob for 50¢). If Jacob were to evaluate each bet by itself, he would pay no more than 9¢ for the first bet and no more than -51¢ for the second bet (i.e. accept no less than 51¢ to take on the burden of the second bet).[17] Thus, if Jacob considers each bet by itself, he will reject both, and thereby forgo a sure gain of 40¢, which he would have if he accepted both offers.

My response to this, as Joyce notes, was to point out that REU-maximizers do not evaluate each bet alone, but instead in the context of the package of bets they are offered. Thus, for example, when Jacob already holds the first bet, the second bet is worth -9¢.[18] Joyce has two criticisms of this response. One is that the reasoning that makes the second bet worth -9¢ when Jacob already holds the first bet is inconsistent with REU theory, and the other is that once Jacob has already *declined* the first bet, there is no reasonable way to make Jacob value the second bet so that he takes it.

Let us begin with the first criticism. Joyce writes:

> "Assume that, for whatever reason, Jacob is at the second shop holding ⟨0.91, –0.09⟩ [i.e. having paid 9¢ for the bet <1, 0>]. [Buchak] tells him to buy ⟨–1, 0⟩ at the price at which he would buy the sum of it and ⟨0.91, –0.09⟩ [i.e. to buy it at -9¢]. This seems to invite Jacob to reason like

[17] Let *i* denote Jacob's total fortune before considering any bets. REU({i+1, 0.3; i+0, 0.7}) = i + $(0.3)^2$((i + 1) – i) = i+0.09, and REU({i-1, 0.3; i+0, 0.7} = i-1 + $(0.7)^2$(i – (i-1)) = i-0.51. Thus, the value that each bet adds to the value of Jacob's current holdings is 9¢ and -51¢, respectively.

[18] The value that the bet adds to the value of Jacob's current holdings is REU({i+1-1, 0.3; i+0+0, 0.7} - REU({i+1, 0.3; i+0, 0.7}) = i - (i + 0.09) = -0.09.

> this: ‘if the bookie offers me ⟨–1, 0⟩ then I can pay her by first *giving* her ⟨0.91, –0.09⟩, which is no skin off my nose since this is a bet that I would not pay anything to buy. I can then sell her ⟨0.09, 0.09⟩ for 9¢ to make up the balance. She gets ⟨1, 0⟩, I spend 9¢, and I dump a bet that has no value to me. So, if she offers more than 9¢ for ⟨–1, 0⟩, I should take it.’
>
> This reasoning is inconsistent with REU. It has Jacob divesting himself of an *asset* and receiving nothing in return. This would only make sense if the asset were *worthless,* by which I man worthless *when sold* not *when bought*. It might seem that ⟨0.91, –0.09⟩ is worthless because Jacob will *buy* it for at most 0¢, but in REU this is not the relevant price when he holds the bet as an *asset.* The bet is valuable to Jacob at the second bookie’s shop because of what he can sell it for, and should be valued accordingly. While giving it away does eliminate the risk in Jacob’s portfolio, it also costs him in expected payoff. When he weighs benefits against costs, Jacob judges that giving ⟨0.91, –0.09⟩ away will cost him 42¢, and so sets that as his selling price.”([14])

The last figure presumably comes from the supposition that selling the asset <0.91, -0.09> is the same as adding <-0.91, 0.09> to Jacob’s fortune when his fortune is <x, x>, and the fact that REU({x - 0.91, 0.3; x + 0.09, 0.7}) = x - 0.42.[19] However, this supposition is false in REU theory *when Jacob already owns the asset*. This is because adding <-0.91, 0.09> to Jacob’s fortune when his fortune is risk-free is different from adding <-0.91, 0.09> to Jacob’s fortune when his fortune has the risk <0.91, -0.09>.[20] The same bet has a different value in different contexts, where context means what the individual already owns. This, indeed, is a lesson of REU theory: different starting assets can alter the risk-profile of a particular bet. So although REU implies that Jacob will set 42¢ as his selling price for <0.91, -0.09> when he does not already own <0.91, -0.09>, it does not imply that Jacob will set 42¢ as his selling price for this asset when he does already own it.

But keep in mind: a given bet has a fixed value in *each* context. The asset <0.91, -0.09> has a fixed value of 9¢ for Jacob in the risk-free context. This means: he indifferent between having the asset (against the background of his risk-free fortune) and having 9¢ (against this same background). He is happy to give his 9¢ in exchange for the asset, and *if he has the asset* he is happy to give it back for 9¢. This is just what it means to value the asset at exactly 9¢. Furthermore, this is true even though if he wouldn’t ‘give

---

[19] REU({i - 0.91, 0.3; i + 0.09, 0.7}) = (i-0.91)+$(0.7)^2$((i+0.09) – (i-0.91)) = i - 0.42.

[20] REU({i+0.91-0.91, 0.3; i-0.09+0.09, 0.7} = REU({i, 0.3; i, 0.7}) = i

the asset away' for 9¢ if he didn't have it already—again, owning an asset puts Jacob in a different position with respect to risk than not owning it.

It is true that in REU, holding fixed a given starting context, a given bet has a different price for an individual when he is buying it (in the sense of paying a fixed amount of money for a receiving a variable amount) and when he is selling it (in the sense of receiving a fixed amount of money for paying a variable amount). This is key to avoiding the Dutch book. But it is false that, again holding fixed a given starting context, a bet that is bought for a given amount of money cannot be sold *back* for that same amount of money.[21] (There are two ways to think of this: one, buying a bet changes the context so that selling it back has the exact same price; two, one's fortune aside from the bet itself is fixed, and the bet considered as a change to this fortune has a fixed indifference price.) Thus, the quoted reasoning really is consistent with REU.

Joyce's other criticism is that Jacob will not, in fact, be in the position of holding <0.91, -0.09> at the second bookie's shop, because he will pass up the first offer. Given this, he will also pass up the second offer. He will thereby pass up a sure gain of 40¢.

Passing up a sure gain would certainly be irrational, so I need to say how Jacob will avoid this, i.e., how he will make a choice other than passing up both bets. In *Risk and Rationality*, I formulated my response to the Dutch book argument in terms of what happens when the individual already holds the first bet,[22] so Joyce is left to extrapolate how I will argue that Jacob will not pass up both bets. He assumes I will say that Jacob will decline the first bet and re-evaluate the price of the second bet:

> "[Buchak] rejects the *Package Principle* (PP), an implicit premise of the Dutch book argument which, in this context, says that Jacob's maximum selling price for <1, 0> at the second bookie's shop should not depend on what happens at the first shop. Buchak claims, to the contrary, that, while Jacob's minimum selling price for <1, 0> should be 51¢ when he has no opportunity to buy it, it should fall to 9¢ once he declines to buy it for 10¢.[23] More generally, this *repricing strategy* says that, upon declining to buy ⟨1, 0⟩ from the first bookie, Jacob should drop his minimum

[21] See Buchak (2013: 205).

[22] Buchak (2013: 211).

[23] Here Joyce cites p. 211 of my book, the discussion of which concerns a case of trying to price the second bet while holding the first bet, rather trying to price the second bet after declining the first bet.

> selling price for the bet from 51¢ to 9¢, which ensures that he will always accept at least one of the two bookies' offers."([9])

He then goes on to show that it would be unreasonable for Jacob to change his selling price for the second bet when his assets haven't changed—i.e. for the mere fact of his declining a bet to make a difference to the value of future bets. I completely agree that this would be unreasonable and ought to be rejected, and Joyce does a good job of bringing out the reasons why. So, while I do want to say that the selling price for the second bet falls from 51¢ to 9¢ when Jacob *accepts* the first bet, I do not want to say that the selling price falls when Jacob *declines* the first bet.

I want to take a different strategy than the one quoted. In particular, Joyce thinks that Jacob must reject the first bet, and then I must come up with some way of explaining how he could accept the second bet (since accepting the second is the only way to avoid forgoing sure gain, given that one has already declined the first). Instead, though, I claim that Jacob will in fact *accept* the first bet—this is what is in fact recommended by REU theory.

The implicit assumption behind the claim that Jacob will reject the first best is that Jacob's attitude towards the first bet is independent of whether he thinks he will be offered the second bet—in other words, that Jacob is *myopic*.[24] An individual is myopic if he makes choices about bets as if each choice is the only choice he will face. Both EU theory and REU theory reject the assumption that individuals are myopic (indeed, they both need to in order to avoid inconsistency), and instead assume that individuals take account of their expectations about the future when deciding what to do. This is justified by noting that the act of taking a bet cannot be fully described without knowing what choices an agent will be offered in the future.

If Jacob is indeed myopic, then he will pass up the first bet, and have no reason to buy the second bet. But if he instead considers the choices he will be offered in the future and chooses in accordance with the plan he thinks best, he will rank his four options from worst to best: buy the first bet but not the second; buy the second bet but not the first; buy neither; and buy both. And since his preferred option will be to buy both,[25] he will formulate and execute the plan of buying both, which has buying the first bet as its

---

[24] See Buchak (2013: 219-220) for a discussion of (and rejection of) myopia.

[25] Just buying the first will be worth -1¢, just buying the second will be worth -1¢, buying neither will be worth 0¢, and buying both will be worth 40¢.

first step. Thus, he will not forgo sure gains, and he has a principled reason for choosing in the way he does.

What will Jacob choose if he doesn't know ahead of time if he will be offered the second bet? That depends: he will have some probability distribution over whether he will be offered a bet in the future, and he will maximize according to this, which means he might accept the bet and might reject it. But regardless, I submit that if we remove Jacob's knowledge of the second bet, then it is no longer the case that forgoing sure gain in this situation will reveal irrationality. The form of the Dutch book argument is that *by the decision-maker's own lights*, he ends up choosing something he disprefers. If Jacob doesn't know ahead of time that he will be offered a second bet, then he is not doing something that is irrational by his own lights. While it is true that he will end up without money he could have had for free, there is no time at which he both recognizes this and repudiates the steps to getting it. Before he is offered the first bet, he does not know that there will be a path to free money if he takes it, and after he declines the first bet, there is no path to free money that he then rejects.

6. Conclusion

The critiques by Thoma and Weisberg and by Joyce have raised important worries for REU theory. These worries have allowed us to get clearer on what the theory is committed to, and indeed what it must be committed in order to avoid the problems raised by these authors. In the case of the Allais paradox, REU theory is committed to the claim that the reasoning that many people cite for the Allais preferences gives rise to the formalism of REU, and furthermore that this reasoning leads to preferences that REU-maximization can capture and EU-maximization cannot. In addition, REU theory is committed to the claim that rank-dependent probability weightings are part of the normative component of decision-making. Finally, REU theory is committed to the claim that when individuals evaluate bets, the price of a given bet depends on which other bets he has and can expect to face. While the challenges raised by the critiques have been important, I remain a committed adherent of REU theory.

Appendix: Mathematica Notebooks

Appendix A: Thoma and Weisberg's original code

I considered only the discrete case, which is given by:

```
umin = −5.;
umax = 5.;
udelta = .1;
pWeight[μ_, σ_, u_]: = 1. −CDF[NormalDistribution[μ, σ], u − .05];

REUA[σ_, x_]: =
    umin + Dot[
        Table[udelta, {i, 100}],
        Table[
          (1.× pWeight[1, σ, u])^x
          , {u, umin + udelta, umax, udelta}]
  ];
REUB[σ_, x_]: =
    umin + Dot[
        Table[udelta, {i, 100}],
        Table[
          (.01 × pWeight[0, σ, u] + .89 × pWeight[1, σ, u] + .1 × pWeight[2, σ, u])^x
          , {u, umin + udelta, umax, udelta}]
  ];
REUC[σ_, x_]: =
    umin + Dot[
        Table[udelta, {i, 100}],
        Table[
          (.9 × pWeight[0, σ, u] + .1 × pWeight[2, σ, u])^x
          , {u, umin + udelta, umax, udelta}]
  ];
REUD[σ_, x_]: =
    umin + Dot[
        Table[udelta, {i, 100}],
        Table[
         (.89 × pWeight[0, σ, u] + .11 × pWeight[1, σ, u])^x
          , {u, umin + udelta, umax, udelta}]
  ];
```

In these formulas, σ represents the variance, and x represents the exponent of the risk function.

Appendix B: Smaller Utility Ratio

I kept the Thoma and Weisberg's definitions of REUA and REUD, but the definitions of REUB and REUC become:

```
REUB[σ_, x_]: =
    umin + Dot[
        Table[udelta, {i, 100}],
```

```
        Table[
          (.01 × pWeight[0, σ, u] + .89 × pWeight[1, σ, u] + .1 × pWeight[1.3, σ, u])^x
          , {u, umin + udelta, umax, udelta}]
  ];
REUC[σ_, x_]: =
    umin + Dot[
        Table[udelta, {i, 100}],
        Table[
          (.9 × pWeight[0, σ, u] + .1 × pWeight[1.3, σ, u])^x
          , {u, umin + udelta, umax, udelta}]
  ];
```

I then verified that the following returned a value of TRUE:

```
REUA[.1, 2.] > REUB[.1, 2.] && REUC[.1, 2.] > REUD[.1, 2.]
```

Appendix C: Skewness

I modified Thoma and Weisberg’s definitions as follows:

```
umin = −5.;
umax = 5.;
udelta = .1;
pWeight[μ_, σ_, a_, u_]: = 1. − CDF[SkewNormalDistribution[μ, σ, a], u − .05];

REUA[σ_, a_, b_, x_]: =
    umin + Dot[
        Table[udelta, {i, 100}],
        Table[
          (1.× pWeight[1, σ, a, u])^x
          , {u, umin + udelta, umax, udelta}]
  ];
REUB[σ_, a_, b_, x_]: =
    umin + Dot[
        Table[udelta, {i, 100}],
        Table[
          (.01 × pWeight[0, σ, b, u] + .89 × pWeight[1, σ, a, u] + .1 × pWeight[1.3, σ, a, u])^x
          , {u, umin + udelta, umax, udelta}]
  ];
REUC[σ_, a_, b_, x_]: =
    umin + Dot[
        Table[udelta, {i, 100}],
        Table[
          (.9 × pWeight[0, σ, b, u] + .1 × pWeight[1.3, σ, a, u])^x
          , {u, umin + udelta, umax, udelta}]
  ];
REUD[σ_, a_, b_, x_]: =
    umin + Dot[
        Table[udelta, {i, 100}],
        Table[
```

$(.89 \times \mathrm{pWeight}[0, \sigma, b, u] + .11 \times \mathrm{pWeight}[1, \sigma, a, u])\hat{}x$
$, \{u, \mathrm{umin} + \mathrm{udelta}, \mathrm{umax}, \mathrm{udelta}\}]$
];

In these formulas, σ represents the variance, *a* represents the skew of $1M and $5M, *b* represents the skew of $0, and *x* represents the exponent of the risk function.

I then verified that the following returned a value of TRUE:

$\mathrm{REUA}[.17, -5,5,2.] > \mathrm{REUB}[.17, -5,5,2.]\ \&\&\ \mathrm{REUC}[.17, -5,5,2.] > \mathrm{REUD}[.17, -5,5,2.]$

Smaller values of σ also returned a value of TRUE, with a variance of ≤ 0.1 compatible with skew values between 0 and -5/5, and a higher variance compatible only with skew values on the increasingly upper end of the range.

Finally, I calculated:
$\mathrm{CDF}[\mathrm{SkewNormalDistribution}[0, .17,5], 1] - \mathrm{CDF}[\mathrm{SkewNormalDistribution}[0, .17,5], .5]$
$\mathrm{CDF}[\mathrm{SkewNormalDistribution}[1, .17, -5], .5] - \mathrm{CDF}[\mathrm{SkewNormalDistribution}[1, .17, -5], 0]$
Both returned a value of 0.00326968. This represents the probability that $0 yields utility between 0.5 and 1, and the probability that $1M yields utility between 0 and 0.5, respectively.

References

Abdellaoui, M, C. Barrios, and P. Wakker (2007), ‘Reconciling introspective utility with revealed preferences: Experimental arguments based on prospect theory’, *Journal of Economics* 138: 356-378.

Briggs, R. (2015), ‘Costs of abandoning the Sure-Thing Principle’, *Canadian Journal of Philosophy* 45(5-6): 827-840.

Buchak, L. (2013), *Risk and Rationality* (Oxford: Oxford University Press).

Buchak, L. (2015), ‘Revisiting Risk and Rationality: a reply to Pettigrew and Briggs’, *Canadian Journal of Philosophy* 45(5-6): 841-862.

Camerer, C.F. (1989), “An Experimental Test of Several Generalized Utility Theories,” *Journal of Risk and Uncertainty* 2: 61-104

Joyce, J. (2017), ‘Commentary on Lara Buchak’s Risk and Rationality’, [this volume]

MacCrimmon, K.R. (1968), ‘Descriptive and Normative Implications of Decision Theory,’ Risk and Uncertainty (eds. Karl Borch and Jan Mossin, New York: St. Martin’s Press), 3-23.

MacCrimmon, K.R. and S. Larsson (1979), ‘Utility Theory: Axioms Versus “Paradoxes”’, *Expected Utility and the Allais Paradox*, eds. M. Allais and O. Hagen (Dordrecht, Holland: D. Reidel).

Machina, M. (1987), ‘Choice under Uncertainty: Problems Solved and Unsolved,’ *Journal of Economic Perspectives* 1(1): 121-154.

Moskowitz, H. (1974), ‘Effects of Problem Representation and Feedback on Rational Behavior in Allais and Morlat-Type Problems,’ Decision Sciences 5: 225-42

Oliver, A. (2003), ‘A quantitative and qualitative test of the Allais paradox using health outcomes,’ *Journal of Economic Psychology* 24: 35-48.

Savage, L. (1954), *Foundations of Statistics* (Dover: John Wiley and Sons).

Schmidt, U. (2004), “Alternatives to Expected Utility: Formal Theories,” Chapter 15 of *Handbook of Utility Theory*, eds. Salvador Barberà, Peter J. Hammond, and Christian Seidl (Boston: Kluwer Academic Publishers): 757-837.

Slovic, P. and A. Tversky (1974), ‘Who Accepts Savage’s Axiom?’, Behavioral Science 19(6): 368-73.

Starmer, C. (2000), ‘Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk,’ *Journal of Economic Literature* 38: 332-382.

Sugden, R. (2004), “Alternatives to Expected Utility: Foundations,” Chapter 14 of *Handbook of Utility Theory*, eds. Salvador Barberà, Peter J. Hammond, and Christian Seidl (Boston: Kluwer Academic Publishers): 685-755

Thoma, J. and J. Weisberg (2017), ‘Risk Writ Large’, [this volume]