

NO ESCAPE FROM ALLAIS: REPLY TO BUCHAK

Johanna Thoma
London School of Economics

Jonathan Weisberg
University of Toronto

ABSTRACT. In *Risk & Rationality*, Buchak (2013) advertises REU theory as able to recover the modal preferences in the Allais paradox. In our (2017) however, we pointed out that REU theory only applies in the “grand world” setting, where it actually struggles with the modal Allais preferences. Buchak (2017) offers two replies. Here we enumerate a variety of technical and philosophical problems with each.

1. INTRODUCTION

In Allais’ (1953) paradox, we face two choices. The first is between options 1A and 1B.

- 1A. (\$1 million, 1)
- 1B. (\$0, .01; \$1 million, .89; \$5 million, .1)

The second choice is between options 2A and 2B.

- 2A. (\$0, .9; \$5 million, .1)
- 2B. (\$0, .89; \$1 million, .11)

Most people prefer 1A over 1B, and 2A over 2B. The prospect of a guaranteed \$1 million is tempting enough to win out against the risky shot at \$5 million.¹ But people are more venturesome when there is no safe option, as in the choice between 2A and 2B.

We are grateful to Lara Buchak and two anonymous referees for feedback on an earlier draft.

¹Throughout, we mean the probabilities to correspond to the agent’s subjective probabilities. Where probabilities are objectively given, e.g., through the experimental setup, we assume agents take these at face value. REU theory takes there to be no essential difference between that case (often referred to as choice under ‘risk’), and the case where probabilities are not externally given (choice under ‘uncertainty’), as long as subjective probabilities can be ascribed to the agent.

Expected utility theory famously forbids this sort of thing. If you're willing to take on an additional .01 risk of empty-handedness in exchange for a .10 chance at \$5 million when there's no safe option, then you should be willing to make the same tradeoff even when there is a safe option.

REU theory is more permissive. Instead of maximizing expected utility, REU theory maximizes *risk-weighted* expected utility:

$$REU(G) = u_1 + \sum_{i=2}^n r(p(u \geq u_{i+1}))(u_{i+1} - u_i).$$

Here G is a gamble with outcomes of utility u_i , ordered from least desirable (u_1) to most desirable (u_n). The value $p(u \geq u_i)$ is the probability of obtaining an outcome with utility at least u_i . And r is a novel component of REU theory, the *risk function*: a non-decreasing map of probabilities into $[0, 1]$, with $r(0) = 0$ and $r(1) = 1$.

The function r captures the agent's attitudes toward risk. If $r(p) < p$ for $p \in (0, 1)$, then the agent is generally risk averse. Buchak's running example of such an r function is $r(p) = p^2$. Given this r function and the utility values

$$u(\$0) = 0,$$

$$u(\$1 \text{ million}) = 1,$$

$$u(\$5 \text{ million}) = 2,$$

we find that $REU(1A) > REU(1B)$ and $REU(2A) > REU(2B)$. So REU theory seems well-placed to capture most people's choices in the Allais problem.

2. GRAND WORLD MODELS

The problem is that no choice is ever really safe or certain. Even if you take the guaranteed \$1 million, your other plans and projects might go well or poorly, superbly or disastrously. Grand world decision problems model this background risk explicitly. Buchak, however, treats the Allais choices as simple problems, in which monetary gains and losses are final outcomes. In any realistic scenario, where there is background risk, this amounts to taking a small world perspective. But because risk-weighted expected utility is non-additive, it has to be applied to grand world problems, as Buchak acknowledges.

We showed that REU theory no longer accommodates the Allais preferences once background risk is modeled explicitly (Thoma & Weisberg, 2017). We constructed a grand world model by replacing each small world outcome with a normal distribution, where the mean matches the small world utility and the height corresponds to its probability. For example, the small world version of gamble 1B is depicted in Figure 1, its grand world counterpart in Figure 2. On this model, REU theory fails to predict

the usual Allais preferences, provided the normal distributions used are minimally spread out.




FIGURE 1. Small world 1B




FIGURE 2. Grand world 1B

Importantly, if we squeeze the normal distributions tight enough, the grand world problem collapses back into the small world problem. Then REU theory can recover the Allais preferences; Buchak's original, small world model would be adequate then.

But the normal distributions have to be squeezed absurdly tight to get this result. A small standard deviation like $\sigma = .1$ lets REU theory recover the Allais preferences.² But it also has outlandish consequences. For example, it entails virtual certainty that a windfall of \$1 million will lead to a better life than the one you'd expect to lead without it. The probability of a life of utility at most 0, despite winning \$1 million, would have to be smaller than 1×10^{-23} .³ Yet the chances are massively greater than that of suffering life-ruining tragedy: illness, financial ruin, etc.

This issue will prove pivotal for Buchak's first reply.

3. BUCHAK'S FIRST REPLY

Buchak's first response is to tweak our grand world model in two ways. First, the utility associated with winning \$5 million is shifted down from 2 to 1.3. Second, all normal distributions are skewed with a shape parameter of 5: positive 5 for the \$0 outcome, negative 5 for the other two. So, for example, the Allais gamble of Figure 2 becomes that in Figure 3.

We'll focus on the second tweak here, the introduction of skew. It rests on a technical error, as we'll show momentarily. But it also wants for motivation.

²Though we need a slightly more severe risk function: $r(p) = p^{2.05}$ instead of $r(p) = p^2$.

³To get this figure we calculate the cumulative density, at zero, of the normal distribution $\mathcal{N}(1, .1)$. In *Mathematica*: `CDF[NormalDistribution[1, .1], 0]`.




FIGURE 3. Skewed grand world 1B

3.1. **Motivational Problems.** Why should the grand world model be skewed? And why in this particular way? Buchak writes:

receiving \$1M makes the worst possibilities much less likely. Receiving \$1M 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 will be fine but not great. Similarly, the utility of \$0 is concentrated around a low 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 great. In other words, \$1M (and \$5M) is a gamble with negative skew, and \$0 is a gamble with positive skew (Buchak, 2017, 2401)

But this passage never actually identifies any asymmetry in the phenomena we're modeling. True enough, "receiving \$1M makes the worst possibilities much less likely." But it also makes the best possibilities much more likely. Likewise, "[r]eceiving \$1M provides security in the sense of making the probability associated with lower utility values smaller and smaller." But \$1 million also makes the probability associated with higher utility values larger. And so on.

The tendencies of large winnings to control bad outcomes and promote good ones was already captured in our original model. A normal distribution centered on utility 1 already admits "some small and diminishing chance that [things] will be fine but not great." It just also admits some small chance that things will be much better than great, since it's symmetric around utility 1. To motivate the skewed model, we would need some reason to think this symmetry should not hold. None has been given.

3.2. **Technical Difficulties.** Setting motivation to one side, there is a technical fault in Buchak’s presentation of the skewed model.

Introducing skew is supposed to make room for a reasonably large standard deviation, while still recovering the modal Allais preferences. Buchak advertises a standard deviation of $\sigma = .17$ for the skewed model, but the true value is actually .106—essentially the same as the .1 value Buchak concedes is implausibly small, and seeks to avoid by introducing skew.⁴




FIGURE 4. The effect of skew on standard deviation

Where does the .17 figure come from then? It’s the scale parameter of the skew normal distribution, often denoted ω . For an ordinary normal distribution, ω famously coincides with the standard deviation σ , and so we write σ for both.⁵ But when we skew a normal distribution, we tighten it; we shrink the standard deviation. Figure 4 illustrates: it shows two normal distributions with the same scale parameter, .17. But the skewed one in yellow is much narrower.⁶

3.3. **Implications.** Of course, what really matters isn’t the value of the standard deviation itself, but the probabilities that result from whatever parameters we choose. And Buchak argues that her model avoids the implausible probabilities we cited in the introduction. How can this be?

Buchak says the skewed model has “more overlap in the utility that \$0 and \$1M might deliver”:

⁴`StandardDeviation[SkewNormalDistribution[1, .17, -5]]`.

⁵Unfortunately, *Mathematica* uses σ for the scale parameter even in skewed normal distributions. This gives the misleading impression that it’s still the standard deviation.

⁶Skewing also shifts the mean, we note.

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. (Buchak, 2017, 2402)

But this overlap is not the problematic quantity we raised in our critique and rehearsed in section 2. The problem was, rather, that a small standard deviation like .1 requires you to think it less than 1×10^{-23} likely that, despite a \$1 million windfall, you will end up with a life no better than what you already expect.

On Buchak's model the corresponding probability is now 3.4×10^{-7} , which is much better, but still absurdly small.⁷ For example, the probability that this paper's second author (Jonathan Weisberg) will die prematurely in the coming year is about 4,000 times greater than that.

Worse yet, the improvement here comes at a steep price: worsened probabilities on the other side. For example, the probability that the life you'll lead with \$1 million will end up at least as good as the one you'd expect with \$5 million is approximately 5×10^{-9} on Buchak's model.⁸ Without the skew Buchak introduces, the equivalent figure would have been a much more reasonable .0013.

4. BUCHAK'S SECOND REPLY

Buchak's second reply is that it wouldn't in fact be a problem if REU theory could only recover the Allais preferences in a "simple" setting. We should think of the Allais problem as a thought experiment: it asks us to abstract away from anything but the immediate rewards mentioned in the problem, and to consider them stand-ins for things that are of ultimate value. For the purposes of this thought experiment, then, the simple decision problem really is grand world: it is a grand world model of a very idealized hypothetical choice situation.

What her simple model shows, according to Buchak, is that REU theory can accommodate people's intuitions regarding such a thought experiment. And this is a success, because this establishes that the theory can accommodate a certain kind of reasoning that we all engage in. Buchak moreover concedes that it may well be

⁷In *Mathematica*, the relevant computation is: `CDF[SkewNormalDistribution[1, .17, -5], .133]`, because .133 is approximately the mean of the skew normal distribution Buchak uses to represent a \$0 outcome (location 0, scale .17, shape 5). Another natural value to consider is the mode, which is approximately 0.063. In that case the relevant probability is even more problematic, approximately 3.5×10^{-8} .

⁸Here we calculate the complement of the cumulative density at 1.17 of the skew normal distribution with location 1, scale .17, and skew -5: `1 - CDF[SkewNormalDistribution[1, .17, -5], 1.17]`, because 1.17 is the mean of the relevant skew normal distribution (location 1.3, scale .17, skew -5).

a mistake for agents to think of the choices they actually face in such simple terms. But she claims this is no problem for her theory.

[I]f 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 REU theory’s ability to capture preferences over highly idealized gambles to which this reasoning is relevant. (Buchak, 2017, 2403)

The first thing to note here is that ordinary agents really do only face “complex” choices as we modeled them. Any reward from an isolated gamble an agent faces in her life really should itself be thought of as a gamble. This is not only true when the potential reward is something like money, which is only a means to something else. Even if the good in question is “ultimate,” it just adds to the larger gamble of the agent’s life she is yet to face. She might win a beautiful holiday in a prize draw, but she will still face 20 micromorts per day for the rest of her life. Even on our deathbeds we are unsure about how a lot of things we care about will play out after we’re gone. REU theory makes this background risk relevant to the evaluation of any individual gamble.

This should make clear just how idealized a thought experiment that is appropriately modeled as a “simple” choice would have to be. It not only has to feature a gamble over ultimate goods, it also has to ask agents to imagine there isn’t and never will be any uncertainty about anything else they value. In standard presentations of the Allais problem, this is at least not something that agents are explicitly asked to do. Rather, the task seems to be the “complex” one of evaluating the prospect of adding a particular sum of money to one’s actual lifetime earnings. And then, as Buchak acknowledges, REU theory fails to rationalize the Allais preferences.

Still, presumably many of us would also display Allais preferences when considering an appropriately idealized, truly “simple” thought experiment. And REU theory can rationalize *that*. But why should we care about accommodating reasoning in such highly idealized decision contexts?

Buchak’s original project was to rationally accommodate the ordinary decision-maker. But now what we are rationally accommodating are at best her responses to thought experiments that are very far removed from her real life. If our model gets things right, then REU theory still has to declare the ordinary decision-maker irrational if she acts in real life as she would in the thought experiment, as presumably

ordinary decision-makers would. And then we haven't done very much to rationally accommodate her. In fact, as far as the project of accommodating ordinary risk aversion is concerned, the burden of proof is still on proponents of REU theory to show that there are any decisions commonly faced by real agents where REU theory comes to a significantly different assessment than expected utility theory. If there are not, then agents may as well just maximize expected utility, as it is easier to compute.

Even if the ordinary decision-maker cannot be accommodated either way, and even if expected utility theory should be used for practical purposes either way, it may of course be that intuitions about highly idealized thought experiments speak in favour of REU theory over expected utility theory as the true theory of rational choice. While we are open to this idea, the proponent of REU theory has a challenge to answer. Why should we trust our intuitive assessments of highly idealized choice situations when we are so prone to misapplying the same kind of reasoning in virtually any other context? Specifically, in the case of the Allais problem, presumably those who report Allais preferences report them whether they are really considering the simple thought experiment, or whether they are really facing the complex choice. What should make us confident in the rationality of their assessment of the thought experiment, when we all agree that their preferences are irrational in the “complex” case?

5. CONCLUSION

Ultimately, if Buchak's first reply fails, and all we can rely on is her second reply, we're left with no reason to abandon expected utility theory as an action-guiding theory in actual choice scenarios. Even if we grant that REU theory is a better theory of rational choice in hypothetical scenarios we never face, this is a much less exciting result than the one *Risk & Rationality* advertised.

REFERENCES

- Allais, Maurice. 1953. Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine. *Econometrica* 21(4): 503–546.
- Buchak, Lara. 2013. *Risk and Rationality*. Oxford: Oxford University Press.
- Buchak, Lara. 2017. Replies to commentators. *Philosophical Studies* 174(9): 2397–2414. doi:10.1007/s11098-017-0907-4.
- Thoma, Johanna & Jonathan Weisberg. 2017. Risk writ large. *Philosophical Studies* 174(9): 2369–2384. doi:10.1007/s11098-017-0916-3.