

RISK WRIT LARGE

Johanna Thoma
University of Toronto

Jonathan Weisberg
University of Toronto

ABSTRACT. Risk-weighted expected utility (REU) theory is motivated by small-world problems like the Allais paradox, but it is a grand-world theory by nature. And, at the grand-world level, its ability to handle the Allais paradox is dubious. The REU model described in *Risk & Rationality* turns out to be risk-seeking rather than risk-averse on one natural way of formulating the Allais gambles in the grand-world context. This result illustrates a general problem with the case for REU theory, we argue. There is a tension between the small-world thinking marshaled against standard expected utility theory, and the grand-world thinking inherent to the risk-weighted alternative.

BUCHAK'S *Risk & Rationality* opens with four examples where the risk-averse choice seems rational, despite violating expected utility theory. These alluring choices appear compatible with Buchak's risk-weighted expected utility theory, however, making it an attractive alternative view of rational choice.

Here we challenge whether REU theory really does accommodate these examples. We will focus on the most famous of the four, the Allais paradox. Our argument is that REU theory struggles to handle this paradox on the theory's own terms. Because REU theory is not partition invariant, it is best understood as a "grand world" theory. It should take into account every possible eventuality of concern to the agent. But the treatment sketched in *Risk & Rationality* follows the usual, "small world" framing appropriate only to partition-invariant theories, like expected utility theory. Moving to the grand-world perspective hampers REU theory's ability to handle the Allais paradox. To recover the usual preferences, strong and implausible assumptions are required.

Thanks to Lara Buchak, Jennifer Carr, Justin Dallmann, Kenny Easwaran, and Sergio Tenenbaum for helpful discussion and feedback.

1. ALLAIS, EU, AND REU

Between the two gambles A and B , which do you prefer?

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

$$B = (\$0, .01; \$1 \text{ million}, .89; \$5 \text{ million}, .1)$$

Most people prefer A to B . Better to walk away with a safe \$1 million than to risk it all for a 10% chance at \$5 million, even if that risk is a meagre 1% chance. Many of these same people prefer C to D given the following choice:

$$C = (\$0, .9; \$5 \text{ million}, .1)$$

$$D = (\$0, .89; \$1 \text{ million}, .11)$$

With a substantial chance of walking away empty-handed already on the table, they are willing to take on an extra 1% risk of empty-handedness in exchange for a 10% chance at \$5 million. But, famously, expected utility theory forbids this combination of preferences (Allais, 1953). If that trade-off is acceptable to you in the second case, it should be acceptable in the first case, too. So you can't simultaneously prefer A to B and C to D .¹

REU theory is more permissive here. It allows us to accept the trade-off between an extra 1% risk of empty-handedness in exchange for a 10% chance at \$5 million in the risky context while rejecting it in the "safe" context, where a guaranteed \$1 million is an option. *Risk & Rationality* illustrates with a simple and plausible model on which the risk-weighted expected utility of A exceeds that of B , yet the risk-weighted expected utility of C still exceeds that of D (*Risk & Rationality*: 71). The model's utility assignments are:

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

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

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

These concave utilities seem plausible enough to us. They don't help expected utility theory explain the usual Allais preferences, though. For that, Buchak argues, we need a new ingredient: the *risk function*.

The risk function alters how probabilities weigh against utilities in a gamble's evaluation. To see how the risk function operates, we start by ordering a gamble

¹For present purposes, we follow Buchak (2013, §4) and bracket the possibility of "redescription". See Pettigrew (2014) for some critical discussion.

from the worst outcome u_1 to the best u_n :

$$G = (u_1, p_1; u_2, p_2; \dots; u_n, p_n).$$

The usual expected utility formula is:

$$EU(G) = u_1 p_1 + u_2 p_2 + \dots + u_n p_n$$

A less familiar but equivalent way of writing this formula weights utility *increases* instead of utilities, as we move from the worst possible outcome to the best:

$$EU(G) = u_1 + \left(\sum_{i=2}^n p_i \right) (u_2 - u_1) + \left(\sum_{i=3}^n p_i \right) (u_3 - u_2) + \dots + p_n (u_n - u_{n-1})$$

The weights here are also different than in the usual expected utility formula: the increase from u_i to u_{i+1} is weighted by the probability that things will be at least as good as u_{i+1} . So we can rewrite this formula:

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

It's these at-least-as-good-as weights that REU theory adjusts using a risk function, r . We apply r to the probability that things will be at least as good as u_{i+1} :

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

If an agent generally gives less weight to the probability that the outcome will be at least as good as u_{i+1} , she will be risk-averse. She will be less influenced by potential gains than a vanilla expected utility maximizer. If instead she gives more weight to these probabilities, she will be risk-seeking:

$$r(p) > p \text{ for all } p \notin \{0, 1\} \Rightarrow \text{risk-seeking}$$

$$r(p) < p \text{ for all } p \notin \{0, 1\} \Rightarrow \text{risk-averse}$$

Risk & Rationality uses $r(p) = p^2$ as its running example of a risk-averse r function. When combined with the u values above, it generates the usual Allais preferences:

$$REU(A) = 1$$

$$REU(B) = .9901$$

$$REU(C) = .02$$

$$REU(D) = .0121$$


So $A > B$ and $C > D$, as desired.

2. A GRAND-WORLD THEORY


Savage (1954) famously noted that every decision really has countless possible outcomes. Even if you take the safe \$1 million, life can still turn out any which way. You might encounter family or health problems that offset the monetary gain, or your winnings might be wiped out in a stock market crash or a lawsuit. Or, things might go the other way, turning out much better than expected, over and above the benefits of your new fortune. So the safe-seeming million is really a gamble, with outcomes of every possible utility.

Expected utility theory can group these numerous possibilities into a handful of “coarse” outcomes because the theory is *partition invariant*, at least when formulated appropriately (Joyce, 1999, 2000). We just need to set the utility of each coarse outcome equal to the weighted average of the numerous, fine-grained eventualities it comprises. Expected utility theory then gives the same results either way. If we calculate the expected utility at the fine-grained level, we get the same evaluation as we do at the coarse-grained level. Expected utility theory gives the same results in the *grand-world* problem as in *small-world* formulations of the same problem.

But REU theory is essentially different in this regard (*Risk & Rationality*: 93). If we lump outcomes together, we alter the gamble’s riskiness. We change its structure, e.g. by making the worst possible outcome more probable, or less bad. Consider a three-outcome gamble with uniform probabilities, and outcomes of utility 0, 1, and 2:



If we lump together the bottom and middle outcomes, and assign the lumped outcome a utility equal to its risk-weighted average, $1/4$, we change the distribution of risk.



The worst outcome isn't quite as bad now, $1/4$ rather than 0 . But it's still not great, and it's now twice as likely you'll end up with that measly $1/4$ of a utile. REU theory is expressly designed to be sensitive to such differences, and the lumping changes its evaluations accordingly: $REU(G) = 5/9$ while $REU(G') = 4/9$.

So REU theory is not partition invariant, but *partition sensitive*. Coarse-graining a gamble's outcomes changes REU theory's recommendations by altering the very risky structure the theory is designed to respond to. For this reason, Buchak says, REU theory must be viewed as a grand-world-only theory. It's to be applied to *final outcomes*: "outcomes whose value to the agent does not depend on any additional assumptions about the world." (*Risk & Rationality*: 93) Using the theory correctly requires fine-graining the outcomes until they specify everything the decision-maker cares about (*Risk & Rationality*: 226–9). Yet we used a small-world rendering of the Allais problem to motivate REU theory in the previous section.

Does it matter?

It does. The model of §1 mishandles the Allais paradox in the grand-world context, at least on one natural way of projecting the small-world Allais gambles onto the big picture. This raises the question whether any plausible model of REU theory can handle the grand-world Allais problem. For if none can, the theory's central motivation is lost.

3. GRAND-WORLD ALLAIS

The safe million of option A is really a gamble. Life might still turn out terrific, terrible, or anywhere in between. How should we represent this gamble?

3.1. Normal Projections. Let's start by considering the status quo. If you're just going about your life as usual, you probably expect things to go reasonably well, though there's a chance they could end up more extreme. You might meet with an unexpected number of life's little setbacks, you might even meet with severe tragedy. On the other hand, things might go significantly better than expected, or even much, much better. How your life will turn out depends on many different events, many flips of fate's coin. So your expectations, we will assume, are captured by the familiar bell-shaped curve of the normal distribution, $\mathcal{N}(\mu, \sigma)$.

Following Buchak, we can set the status quo as the zero-point of our utility scale. So before the Allais gambles come into the picture, your expectations are normally distributed around the mean $\mu = 0$.

What should the standard deviation σ be? We will start with the somewhat arbitrary but charitable assumption that $\sigma = .2$. Smaller values of σ are better for REU theory, as we'll see, and $\sigma = .2$ is quite small. On Buchak's utility scale, a gain of \$1 million increases your utility from 0 to 1, which is five standard deviations if $\sigma = .2$. That means $\sigma = .2$ is so small, you are more than .9999997 confident that life without the \$1 million will be less good than what you would normally expect with the \$1 million.

Grand-world versions of the Allais gambles can now be obtained by adjusting your expectations from the status quo. For example, the "safe" \$1 million of gamble *A* shifts the mean up to $\mu = 1$. If you gain a million dollars right now, other events in your life could still turn out any which way. But most likely, things will go as expected, with the \$1 million improving things in the way one ordinarily hopes. In other words, gamble *A* corresponds to the normal distribution $\mathcal{N}(1, .2)$ depicted in Figure 1.

What about gamble *B*? It has three small-world outcomes: \$0, \$1 million, and \$5 million. So we replace each of these with a normal distribution centered on its utility, though scaled down according to its probability. Applying the same method to gambles *C* and *D* we get the distributions illustrated in Figure 1.

These are continuous distributions, whereas Buchak defends REU theory in a finite, discrete setting. But we can bridge the gap in a couple of ways, and it turns out not to matter which we choose. So we reserve discussion of this wrinkle for the Appendix, and proceed with our continuous approach.

3.2. The Challenge for REU Theory. What does REU theory say about our grand-world Allais gambles? Assuming $r(p) = p^2$, we find that REU theory is actually risk-seeking! *B* is now preferable to *A*:

$$REU(A) \approx .887$$

$$REU(B) \approx .900$$

While *C* continues to be preferable to *D*:

$$REU(C) \approx -.073$$

$$REU(D) \approx -.079$$

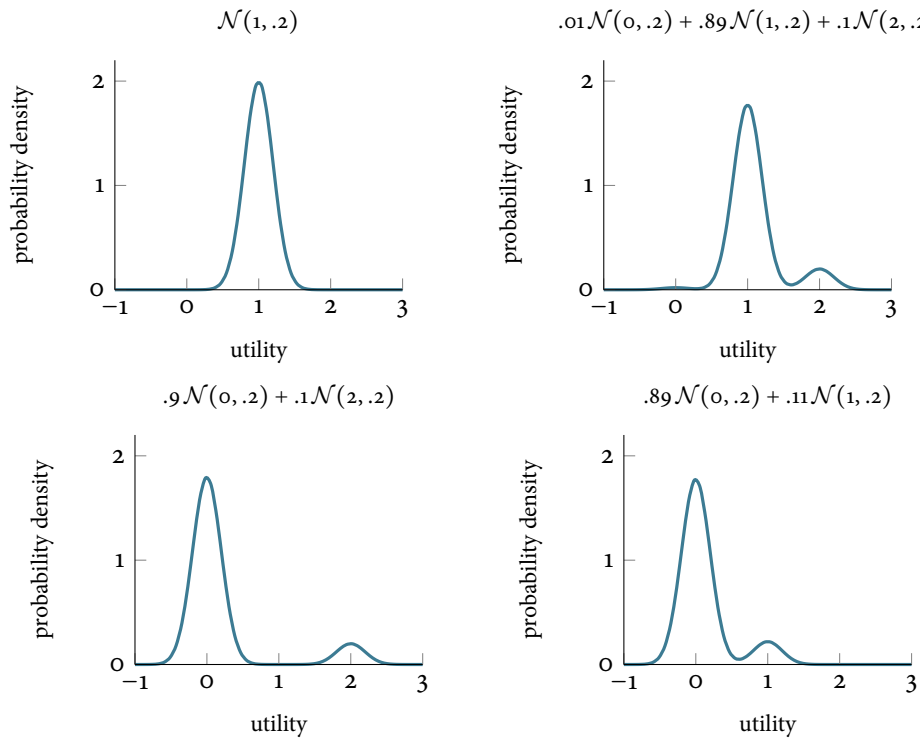


FIGURE 1. Grand-world gambles *A* and *B* (top), and *C* and *D* (bottom)

In other words, REU theory now apes EU theory's preferences. It does no better at explaining risk-aversion in the Allais paradox than the theory it was meant to replace.

In a way, this is not surprising. From the grand-world perspective, option *A* is not really a sure thing. Though it may secure the \$1 million, it cannot secure a life of utility 1 rather than utility 0. From the grand-world perspective, *A* is risky, much like *B*.

So the challenge for REU theory is a tension between the grand-world context it embraces, and the small-world thinking it seeks to validate. The appeal of option *A* is its certainty, the opportunity to avoid any risk. But REU theory insists on the grand-world context, where that appeal dissolves. At the grand-world level, *A* is risky and may not be preferable as a result. Indeed, we've just seen that it is not preferable on one natural way of modeling the grand-world context.

Other models might be more friendly to the REU theorist's cause, though. To meet the challenge, they might try altering some of the parameters we've introduced. Or they might defend an altogether different model.

We have examined several variations on the present model and found them all wanting. To summarize, the only way we have found for REU theory to recover the Allais pattern is to make σ implausibly small, and the risk function r implausibly extreme or specific. We go into more detail in the next section, and then discuss the broader significance of our findings in §5.

4. TWEAKING THE MODEL

4.1. Varying Levels of Risk Aversion. Would a more risk-averse r function make A preferable? Let's consider $r(p) = p^x$ with other values of x besides 2.

We examined values of x ranging from 1 to 10 at intervals of .01.² The result: A did eventually become preferable to B , as expected. But D became preferable to C first. As described in the Appendix, a careful search suggests strongly that there is no value of x that recovers the Allais pattern, given $\sigma = .2$. Evidently, $REU(D) > REU(C)$ for $x \geq 2.57$, while $REU(A) > REU(B)$ only for values of x strictly greater than 2.57.

But even if there were a value of $x > 2$ that succeeded, we would have reservations. The risk function $r(p) = p^2$ is already pretty extreme. An agent with this risk function would reject a gamble that gives her a 50% chance of losing \$100 and winning \$299, given utility linear in dollars. At $r(p) = p^3$, she would reject such a gamble up until a potential gain of \$699.³

4.2. Larger Values of σ . We suggested that $\sigma = .2$ is implausibly small when we first introduced this value, adopting it only to be charitable. If we make the model more realistic by increasing σ , the Allais pattern fails to emerge, as expected. Instead, REU theory's preference for B over A only gets stronger. We explored the range $.2 \leq \sigma \leq 1$ at .01 intervals and found that $REU(A) - REU(B)$ only becomes more negative as σ increases. (Again, see the Appendix for details.)

4.3. Smaller Values of σ . To put our cards on the table though, we selected $\sigma = .2$ as our working "small" value because it's about as small as σ can get before the above results fail to hold. If we set the standard deviation lower, the Allais pattern can be recovered.

At $\sigma = .1$, a slight adjustment to the r function is all we need to recover the pattern. Just bump x from 2 up to 2.05 and we get the desired result. We have already expressed reservations about $r(p) = p^2$ implying an extreme level of risk-aversion. All the more for $r(p) = p^{2.05}$. But bracket that concern for a moment.

²Why start with 1 instead of 2? Just to be thorough.

³These are small-world examples, but with an appropriate back story, the small-world problem can be made the same as the grand-world problem. For example, the gambles might be offered by God on the last day of your life, with currency replaced by heavenly utilities.

of accounting for Joe Average's risk-aversion. Indeed, as we have seen, Joe Average can become Joe Impossible quite easily, even while allowing that a "safe" \$1 million really is quite safe ($\sigma = .2$).

5. DISCUSSION

Stepping back, a larger point emerges. There is a kind of paradoxical irony to REU theory.

The theory is meant to sympathize with our aversion to uncertainty. It allows us to eschew options whose outcomes are less predictable—more "spread out" as Buchak says—in favour of options whose outcomes are more determinate. To achieve this effect though, the theory appears to bind itself to the grand-world problem. It rejects the additive approach of expected utility theory, apparently sacrificing the ability to work at the small-world level as a result.

The irony is that, at the grand-world level, everything is spread out. Every choice has innumerable possible outcomes, and it is never certain how one's choice right now will play out in the grand scheme of things. Even a "safe" \$1 million might leave you destitute and miserable in the end. And that possibility threatens to undercut the initial motivation for the theory. It doesn't necessarily recommend the "safe" option anymore once it's in scare-quotes, which it is in the grand scheme of things.

REU theorists might try to answer this challenge a number of ways. Let's explore two of them, and see what challenges they face.

5.1. Response #1: Small Worlds After All. REU theorists might point out that people don't usually think about the Allais gambles anything like the way we have described them. One doesn't normally view them from the grand-world perspective, but rather just sees the \$1 million as a guaranteed improvement by 1 utile over the status quo. And, framed this way, REU theory easily sympathizes, as Buchak's original model shows.

But this may be sympathy for the devil. Perhaps it's a descriptive truth that people view these gambles in small-world terms. But as we have seen, Buchak herself claims that REU theory forbids small-world thinking, because the theory is partition-sensitive.

Could framing a decision problem in small-world terms be permissible, despite REU theory's partition-sensitivity? Given partition-sensitivity, using a more fine-grained description of a decision problem can change the theory's recommendations. It is usually held that we should go with the most fine-grained description in such circumstances. Two main considerations support this view.

First, we may think that the grand-world decision problem is ultimately what we should be solving. If small-world decision problems are just attempts at modeling the grand-world decision problem, then partition-sensitivity implies that small-world problems can be bad models. One reason for thinking it's ultimately the grand-world decision problem we ought to be solving is that fine-grained outcomes are the location of value. And decision theory is supposed to capture how to best achieve ends we value.

Second, one might think that a description of a decision problem should capture everything that is *relevant* for the agent. One criterion for relevance could be that any detail that may change the agent's decision should be included in the specification of the decision problem. And then, under partition-sensitivity, small-world decision problems may leave out relevant detail.

There may be some room for challenging these ideas. We could be permissive about the framing of decision-problems despite partition-sensitivity. In response to the arguments just provided, one might hold that the agent herself can decide how much detail is relevant to her decision, and that value resides at whatever level of description she chooses. McClennen expresses this view when he writes, "If the world in fact opens to endless possibilities, still evaluation of risks and uncertainties requires some sort of closure [...] Wherever the agent sets his horizons, it is here that he will have to mark outcomes as terminal outcomes—as having values that may be realized by deliberate choice, but nevertheless as black boxes whose contents, being undescribed, are evaluatively irrelevant." (McClennen, 1990, 249)

The problem with this response is that it makes the recommendations of REU theory highly sensitive to framing and context. The detail that the small-world version of the Allais problem leaves out is detail that the agent will likely find relevant in other choice contexts. In many contexts, it will presumably matter, for instance, how much interest you can get on your \$1 million, or whether some disease will keep you from enjoying it. If we are permissive about framing, then we need to hold that the choice context changes what the agent finds relevant, and at what level of description she assigns ultimate value. We suspect most decision theorists will not be willing to bite this bullet.

Of course, even if ideal REU theory is a grand-world theory, bounded rationality might still call for smaller, more tractable frames. But then small-world EU would be the better heuristic. As Buchak notes, given some plausible assumptions, EU maximizers and REU maximizers make pretty much the same choices in the grand-world problem (*Risk & Rationality*: 227–8). And since EU is partition-invariant, small-world EU will match grand-world EU, which will closely match grand-world

REU. Whereas small-world and grand-world REU can come well apart, as we've seen.

So the challenges here are, first, to rationalize the use of small-world reasoning. And then, second, to rationalize the use of REU-maximization rather than EU-maximization in small-world problems.

5.2. Response #2: Varying the Variance. Gamble *A* doesn't just promise safety in that you definitely get the \$1 million. It also promises the safety of financial security. Those with \$1 million in the bank are less vulnerable to many of life's setbacks; they are not as easily ruined as the rest of us. And this points up an unrealistic feature of our model: we assumed the same value of σ across the board. But really, it should vary.

If \$1 million shrinks σ by providing a measure of financial security, then \$5 million shrinks it even more. So we need three values to replace σ , one for each small-world outcome: σ_0 for utility 0, σ_1 for utility 1, and σ_2 for utility 2. The general constraint we have to work with is:

$$\sigma_0 > \sigma_1 > \sigma_2$$

For example, we might set $\sigma_0 = .4$, $\sigma_1 = .2$, and $\sigma_2 = .1$. If we do, we find the same problem as before. REU theory is still risk-seeking given $r(p) = p^2$, preferring *B* over *A* and *C* over *D*.⁴

We conducted a search of what values of x and of the triple $(\sigma_0, \sigma_1, \sigma_2)$ may recover the Allais preferences. We let the σ_i range from .1 to 1, and let x range from 1 to 4 at intervals of 0.1, taking values outside of these ranges to be too implausible to consider. Some values within these ranges indeed recover the Allais preferences. However, as before, a highly specific combination of values is needed. Solutions only appear to exist for $2.1 \leq x \leq 2.7$. So recovering the Allais preferences requires a fairly extreme level of risk aversion. And even then, very specific combinations of σ_i are needed, involving very small σ_1 and σ_2 . The Appendix links to a *Mathematica* notebook containing more details on the search and our results.

We thus face the same problems as before. The Allais preferences cannot be recovered with a large and robust range of REU models. And the REU models that do recover them involve implausibly high levels of risk aversion and security.

Even if more plausible models could be found, there might still be cause for concern. Though it's true that *A* is safe in the sense of promising financial security, its advantage over *B* is supposed to come from the other kind of safety it promises: the short-run guarantee of walking away \$1 million dollars richer. So there is a concern

⁴ $REU(A) \approx 0.944$, $REU(B) \approx 0.945$, $REU(C) \approx -0.0263$, $REU(D) \approx -0.0333$.

about getting the right results in the wrong way. If REU theory does manage to capture the usual Allais preferences, but only by appealing to reasons separate from those that drive actual agents to have those preferences, it becomes doubtful whether capturing those preferences really vindicates the theory.

6. CONCLUSION

The moral we draw is that REU theory doesn't clearly handle the very problems it was designed to solve. It's not that REU theory is flat-out inconsistent with the usual Allais preferences in the grand-world context. To the contrary, we provided some grand-world REU models that suggest the opposite. The trouble is that the only such models we found weren't very plausible. They come too close to the small-world problem by setting σ implausibly low.

Of course, there are many other shapes the risk function might take besides $r(p) = p^x$, and other shapes may do better. Also, there are surely more realistic ways of projecting the Allais gambles onto the grand-world context. We only scratched the surface on one of these, when we briefly considered using different σ_i 's for different small-world outcomes. So there may yet be models of REU theory that fit the bill. For the theory to live up to its promise, however, we need to actually identify plausible candidates. Until we do, it's unclear how successful REU theory really is at achieving its own ends.

7. APPENDIX

Here we describe the results of §§3–4 and how they are obtained. The *Mathematica* code described below is available for download at the following URL:

<http://www.utm.utoronto.ca/~weisber3/articles/REU.nb>

Code for the three- σ model discussed in §5.2 can also be found there.

7.1. Analogue vs. Digital. As noted in §3.1, Buchak defends REU theory in a discrete, finite context, though we used continuous distributions in the main text. There are two ways to bridge this gap.

The first way is to extend REU theory to the continuous context. Given a gamble represented by a continuous density $p(u)$ over an interval of utilities $[u_{\min}, u_{\max}]$, we calculate its *REU* using p 's cumulative distribution function, $P(u)$:

$$REU(G) = u_{\min} + \int_{u_{\min}}^{u_{\max}} r(1 - P(u)) du$$

This formula is just what we get from the discrete *REU* formula if we regard G as a discrete gamble with evenly spaced outcomes, with utilities Δ apart, and then let $\Delta \rightarrow 0$.

The second way is to work with discrete gambles directly. For example, we could let Allais' gamble A have finitely many outcomes, with utilities:

$$-5.0, -4.9, -4.8, \dots, 4.8, 4.9, 5.0$$

We could then assign discrete probabilities that approximate the continuous distribution $\mathcal{N}(1, .2)$. For example, an outcome of utility 1 would be assigned:

$$p(u = 1) = \int_{.95}^{1.05} \mathcal{N}(1, .2) du$$

And likewise for the other 100 possible utility values.

These “fragmentations” of the continuous gambles give us finite and discrete representations of the grand-world context, as illustrated in Figure 2. We can then

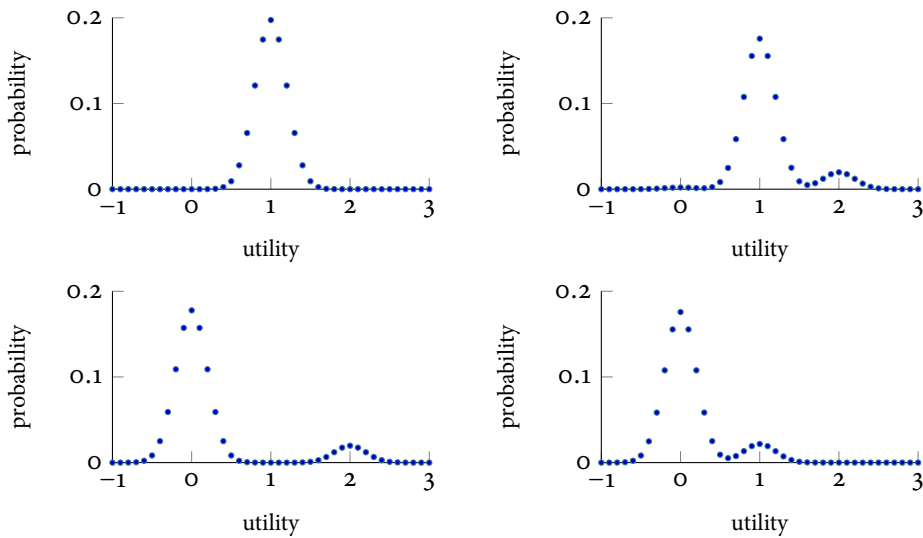


FIGURE 2. Fragmented gambles A (top-left) through D (bottom-right)

apply the standard, discrete *REU* formula Buchak defends.

Whether we go discrete or continuous, we will have to work with numerical approximations. The normal distribution is central to both approaches. But, if we take the discrete approach yet make it very fine, the numerical *REU* values can be arbitrarily close to those of the continuous approach. In fact, the discrete approach needn't be very fine at all to match the results of the continuous approach that we will use here. The discrete model just described, where possible utilities range from

-5 to 5 at intervals of 0.1 , gives the same results as the continuous model. So we will describe only the continuous approach here. (The *Mathematica* notebook linked above provides code for both approaches.)

7.2. Programming the Model. We start by defining the Q -function, the complement of the normal distribution's cumulative distribution function:

$$Q[\mu_, \sigma_] := 1. - \text{CDF}[\text{NormalDistribution}[\mu, \sigma], u]$$

Then we define four functions, one for each Allais gamble, to compute its REU given a standard deviation σ and a power x for the risk function $r(p) = p^x$:

$$\begin{aligned} REUA[\sigma_, x_] &:= -5. + \int_{-5.}^{5.} (1. \times Q[1., \sigma])^x du // N \\ REUB[\sigma_, x_] &:= -5. + \int_{-5.}^{5.} (.01 \times Q[0., \sigma] + .89 \times Q[1., \sigma] + .1 \times Q[2., \sigma])^x du // N \\ REUC[\sigma_, x_] &:= -5. + \int_{-5.}^{5.} (.9 \times Q[0., \sigma] + .1 \times Q[2., \sigma])^x du // N \\ REUD[\sigma_, x_] &:= -5. + \int_{-5.}^{5.} (.89 \times Q[0., \sigma] + .11 \times Q[1., \sigma])^x du // N \end{aligned}$$

We have chosen -5 and 5 as the minimum and maximum utility values because they are quite extreme, and more extreme values are so improbable as to have no impact on the results that follow.

7.3. Results for $\sigma = .2$. First we verify that REU theory is risk-seeking in the grand-world context given $r(p) = p^2$ and $\sigma = .2$:

$$\begin{aligned} REUA[.2, 2.] < REUB[.2, 2.] \ \&\& \ REUC[.2, 2.] > REUD[.2, 2.] \\ \text{True} \end{aligned}$$

So we consider other risk functions of the form $r(p) = p^x$, and examine the range $1 \leq x \leq 10$ at intervals of $.01$:

$$\begin{aligned} \text{AminusB} &= \text{Table}[\{x, REUA[.2, x] - REUB[.2, x]\}, \{x, 1., 10., .01\}]; \\ \text{CminusD} &= \text{Table}[\{x, REUC[.2, x] - REUD[.2, x]\}, \{x, 1., 10., .01\}]; \\ \text{ListPlot} &[\text{AminusB}, \text{CminusD}] \end{aligned}$$

Though it's not immediately obvious from the graph (Figure 3), $REU(D)$ overtakes $REU(C)$ before $REU(A)$ overtakes $REU(B)$. By the time $x = 2.57$, C is no longer

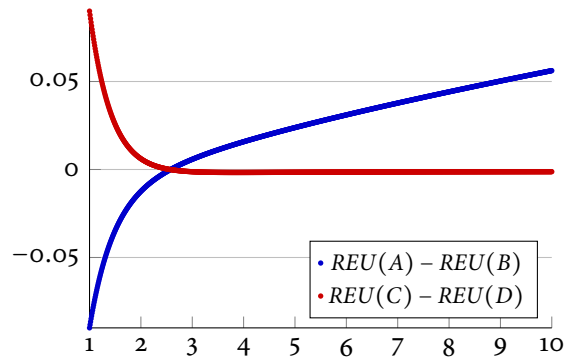


FIGURE 3. $REU(A) - REU(B)$ and $REU(C) - REU(D)$ for $2 \leq x \leq 3$

preferable to D while A has yet to become preferable to B :

```
REUA[.2, 2.57] < REUB[.2, 2.57] && REUC[.2, 2.57] < REUD[.2, 2.57]
True
```

Evidently, there is no r function of the form $r(p) = p^x$ capable of producing the Allais pattern in the grand-world context when $\sigma = .2$.

7.4. **Results for $\sigma > .2$.** Unsurprisingly, increasing σ isn't promising. Still, for completeness, we check the range $.2 \leq \sigma \leq 1$ at $.01$ intervals to see if the Allais pattern might re-emerge:

```
AminusB = Table[{σ, REUA[σ, 2.] - REUB[σ, 2.]}, {σ, .2, 1., .01}];
ListPlot[AminusB]
```

As expected, increasing σ only decreases the appeal of A relative to B : see Figure 4. So we turn instead to examine smaller values of σ .




FIGURE 4. $REU(A) - REU(B)$ for $.2 \leq \sigma \leq 1$

7.5. **Results for $\sigma < .2$.** Setting $\sigma = .1$ doesn't by itself recover the Allais pattern. Given $r(p) = p^2$, we still have $REU(A) < REU(B)$:

$$REUA[.1, 2.] < REUB[.1, 2.]$$

True

But a slight increase in risk-aversion is now sufficient to reverse this. Just by setting $r(p) = p^{2.05}$, we recover the Allais pattern:

$$REUA[.1, 2.05] > REUB[.1, 2.05] \ \&\& \ REUC[.1, 2.05] > REUD[.1, 2.05]$$

True

On the other hand, too large an increase revives the problem of D being preferable to C . For example, if we set $r(p) = p^{2.9}$:

$$REUC[.1, 2.9] < REUD[.1, 2.9]$$

True

And as σ increases, the range of viable x -values narrows. For example, at $\sigma = .15$, x has to be in a subinterval of $(2.2, 2.7)$:

$$\{REUA[.15, 2.2] > REUB[.15, 2.2], REUC[.15, 2.2] > REUD[.15, 2.2]\}$$

{False, True}

$$\{REUA[.15, 2.7] > REUB[.15, 2.7], REUC[.15, 2.7] > REUD[.15, 2.7]\}$$

{True, False}

And at $\sigma = .19$, x has to be in a subinterval of $(2.5, 2.6)$:

$$\{REUA[.19, 2.5] > REUB[.19, 2.5], REUC[.19, 2.5] > REUD[.19, 2.5]\}$$

{False, True}

$$\{REUA[.19, 2.6] > REUB[.19, 2.6], REUC[.19, 2.6] > REUD[.19, 2.6]\}$$

{True, False}

So increasing σ tightly constrains x .

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.

- Huck, S. & W. Muller. 2012. Allais for All: Revisting the Paradox in a Large Representative Sample. *Journal of Risk and Uncertainty* 44(3): 261–293.
- Joyce, James M. 1999. *The Foundations of Causal Decision Theory*. New York, NY: Cambridge University Press.
- Joyce, James M. 2000. Why We Still Need the Logic of Decision. *Philosophy of Science* 67(S1): S1–S13.
- McClennen, Edward F. 1990. *Rationality and dynamic choice: Foundational explorations*. New York, NY: Cambridge University Press.
- Pettigrew, Richard. 2014. Buchak on risk and rationality III: The redescription strategy. <http://m-phi.blogspot.ca/2014/04/buchak-on-risk-and-rationality-iii.html>.
- Savage, Leonard J. 1954. *The Foundations of Statistics*. Hoboken, NJ: Wiley.