Can risk aversion survive the long run?*

Hayden Wilkinson

Forthcoming in *Philosophical Quarterly*

Abstract

Can it be rational to be risk-averse? It seems plausible that the answer is yes—that normative decision theory should accommodate risk aversion. But there is a seemingly compelling class of arguments against our most promising methods of doing so. These *long-run* arguments point out that, in practice, each decision an agent makes is just one in a very long sequence of such decisions. Given this form of dynamic choice situation, and the (Strong) Law of Large Numbers, they conclude that those theories which accommodate risk aversion end up delivering the same verdicts as risk-neutral theories in nearly all practical cases. If so, why not just accept a simpler, risk-neutral theory? The resulting practical verdicts seem to be much the same. In this paper, I show that these arguments do not in fact condemn those risk-aversion-accommodating theories. Risk aversion can indeed survive the long run.

Keywords: *risk aversion; risk-weighted expected utility theory; expected utility theory; long-run arguments; dynamic choice.*

Word count: 8,977

*I am grateful to Teru Thomas, Brian Hedden, Timothy L. Williamson, Alan Hájek, and Riley Harris for their detailed comments on drafts of this paper, and to Christian Tarsney and a reading group at the Australian National University for helpful discussions. I also thank the Forethought Foundation for Global Priorities Research for their generous financial support.

1 Introduction

Consider two options for how to spend $50. You could donate it to a charity that performs surgeries to treat trachoma, an infection of the eye which causes blindness, and which is common among the world's poorest. Donate to that charity and one additional person will be cured of blindness and, we can suppose, is sure to have a better life. Alternatively, you could donate your $50 to a charity that lobbies wealthy governments to slightly increase their funding of clean water and sanitation programmes in low-income countries. Do so and you have some probability[1], perhaps one in 100, of making the difference in persuading such a government, causing an influx of funding, improving sanitation in poor countries, and thereby preventing far more cases of trachoma (for which poor sanitation is a key risk factor). But, more likely, it will make no difference at all to those who suffer from trachoma. In a useful simplification of reality, we might suppose that your options are as follows.

A : 1 person cured of blindness with probability 1.

B: 101 people cured of blindness with probability $1/100$; 0 people cured otherwise

Which of these options is better, instrumentally, given your uncertainty? (I will focus on *moral* betterness here and throughout, but analogous cases and arguments arise for prudential and other forms of betterness.)[2] One common approach is *expected value theory*, which says that an option L is at least as good (instrumentally) as another option M if and only if L brings about at least as great an expectation of (moral) value as M. This expected value (or EV) is the sum of the value of the option's possible outcomes, each weighted by its probability of eventuating. In the comparison above, the (moral) value of each outcome can be treated as directly proportional to the number of people cured of blindness, because this is the only difference between outcomes, because the ordering will be determined entirely by those differences, and because, morally, it seems overwhelmingly plausible that it is equally valuable to cure each additional person of blindness regardless of how many others have been cured—the good of curing people of blindness does not have diminishing marginal value (assuming that those people do not interact). So, the expected value of A will be 1, and the expected value of B will be $101 \times 1/100 = 1.001$, which is greater. Expected value theory says that it is better to take the gamble to cure 101 people at once.

But you might find it more intuitive that A is better. After all, it will almost certainly turn out better than B—your donation to trachoma surgery is guaranteed to benefit someone, while your donation towards lobbying wealthy governments will almost certainly go to waste. You might think that risky options like this are not as good as their expected value suggests—that rational decision-making should accommodate a greater degree of *risk aversion*. But should it?

According to *long-run* arguments, risk aversion should not be accommodated.[3] Why not? In practice, agents face similar decisions many times throughout their lives. So, how we rationally ought to compare

[1] This probability, and those used throughout, may be subjective degrees of belief or evidential probabilities, whichever kind you think relevant for normative decision theory. For the purposes of this paper I will remain agnostic on which kind that is. I will also assume, for simplicity, that outcomes can always be assigned precise probabilities, *contra* Knight (1921) and others.

[2] To translate the discussion into prudential (or any given form of) betterness, we need a privileged cardinal measure of how good each outcome is: an objective value function. If no such value function exists, expected value theory makes little sense. Risk aversion may not make sense either if it is modelled only using expected utility theory (see below).

[3] Such arguments are presented, sometimes against particular theories of risk aversion and sometimes against accommo-

our options depends not on what our theory says about how to compare options in single decisions in isolation, but on what it says about how to compare options in *dynamic* decision situations. By plausible dynamic choice rules[4], over the long run of an agent's life, theories that accommodate risk aversion are alleged to give surprising verdicts: precisely the same verdicts as expected value theory, in nearly all practical cases. Such theories "...[fail] to offer a true alternative to expected [value] theory" (Thoma, 2019, p. 230).[5] If those theories simply parrot the verdicts of expected value theory in practice, or at least do so in almost all of the cases we most care about, we seem to have little reason to adopt them at all.

Such arguments have problems. For one, even the most compelling such arguments seem to depend on a somewhat controversial principle of dynamic decision-making: resolute choice.

For another, more troubling problem, there turn out to be many exceptions (including many exceptions in practice) to the claim that those theories merely parrot expected value theory. When long-run arguments are stated informally, it may indeed seem plausible that those theories must come to always agree. But, when we attempt to rigorously prove the necessary results, there emerge unavoidable exceptions. As I will demonstrate, particular risk-aversion-accommodating theories often fail to agree with expected value theory—namely, when our options have equal expected value, when we face decisions with particularly high stakes, when we adopt certain risk-averse attitudes, and when we make decisions near the ends of our lives. So these theories do not *always* agree with expected value theory; not even close. These are sufficiently broad and important classes of cases that disagreement therein is enough to distinguish those theories from expected value theory.

So, contrary to the conclusions drawn by others, risk-aversion-accommodating theories *can* offer a genuine alternative to expected value theory. We need not banish risk aversion from decision-making just yet.

2 Models of risk aversion

How, exactly, can a theory of instrumental betterness accommodate risk aversion? Throughout this paper, I will consider the two most common (and, to me, most plausible) proposals.

The first such proposal is *expected utility theory*. It says that an option L is at least as good as another option M, if and only if L brings about at least great an expectation of *utility* as M. This may sound a lot like expected *value* theory from earlier. But here I mean *utility* in a decision-theoretic sense, meaning something quite different from (moral) *value*—it also represents the agent's attitude to risk.[6] Depending on how strongly the agent prefers obtaining good outcomes with high probability, the utility of an outcome

dating risk aversion in general, by, e.g. Thoma (2019), Stefánsson (2020), Zhao (2021), Briggs (2019), Hájek (2021, §.2), and Baron (2000, p. 244).

[4]See Thoma (2019, p. 241) and Zhao (2021, p. 17).

[5]Thoma's criticism is aimed at Lara Buchak's risk-weighted expected utility theory but, as I'll show, an analogous argument can be made against versions of expected utility theory that allow risk aversion.

[6]I will assume throughout, as Buchak (2013) does, that the normatively relevant risk attitude is one held by the agent making the decision. But we might adopt a version of expected utility theory, or risk-weighted expected utility theory below, that isn't agent-relative—there might be some objectively correct risk attitude that happens to be a risk-averse one, and/or evaluations might need to be agreed upon by a range of risk attitudes. The discussion throughout will apply to these alternative versions too.

can be *any* increasing function of its value—one outcome will still have greater value than another if and only if it has greater utility, but *how much greater* its value is can vary.

Expected utility theory allows us to accommodate general risk aversion with respect to (moral) value by allowing the agent's utility function to be *concave*: the greater the value of an outcome, the less difference to its utility is made by some additional value. Such a utility function looks like this:

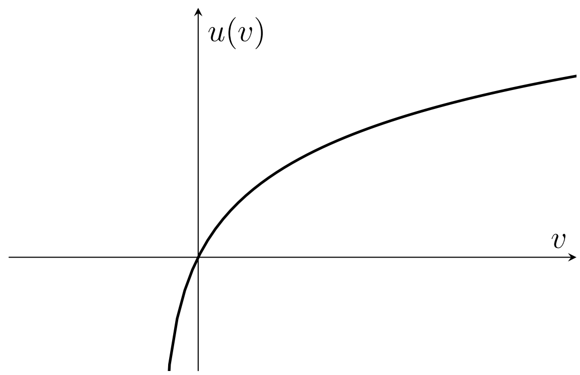

Figure 1: A utility function $u(v)$ that is concave with respect to (moral) value.

If the agent making the decision has such a utility function (with sufficient concavity), expected utility theory can deliver the risk-averse verdict in the case from earlier. Recall that A is the option of curing one person's blindness with certainty, and B the option of curing the blindness of 101 people with a probability of one in 100. In terms of *utility*, the options might be represented as follows.

A : (an outcome of) utility 1 with probability 1.

B: utility 50 with probability $1/100$; utility 0 otherwise

We would then assign an expected utility 1 to the safe option, A, and expected utility 0.5 to B. So, it is indeed better to save one person for sure than to gamble on saving 101 lives, at least on this risk attitude. Thus, expected utility theory can vindicate the risk-averse verdict we might want here.

An alternative and widely-discussed proposal for accommodating risk aversion is Lara Buchak's *risk-weighted expected utility theory*.[7] According to this theory, an option L is at least as good as another option M if and only if the risk-weighted expected utility (REU) of L is at least as great as that of M. This REU is calculated as follows.

$$\text{REU}(L, r) = u_1 + \sum_{j=2}^{n} (u_j - u_{j-1}) r\Big(P(L \geq u_j)\Big)$$

Here, $u_1, u_2, u_3, ..., u_n$ are the utilities of the possible outcomes of L, ordered from lowest to highest. (I will assume that these utilities are given by a utility function that is *linear* with respect to value, both for simplicity and to highlight the differences between the two theories.)[8] The REU formula asks us to start

[7] Its promise comes in large part from its ability to accommodate the intuitively rational preferences described by Allais, which expected utility theory cannot accommodate—see Buchak (2013, pp. 31-4).

[8] The results below concerning REU theory—Theorems 2 and 4—apply also for REU theory with non-linear utility functions.

with the utility of the worst possible outcome, given by u_1, and sum the amounts by which L could exceed u_1 (and u_2, and so on). But we weight those amounts, not by the *probability* that we exceed the previous u_{j-1} by that amount, but by some function *of* that probability.

This risk function, r, represents the agent's risk attitude. By definition, r must be real-valued, non-decreasing, and satisfy $r(0) = 0$ and $r(1) = 1$ (**?**, p. 49). And it allows us to accommodate general risk aversion with respect to (moral) value by allowing the agent's r function to be *convex*, and to take values of $r(p)$ less than p for all probabilities between 0 and 1. Such an r function looks like this:

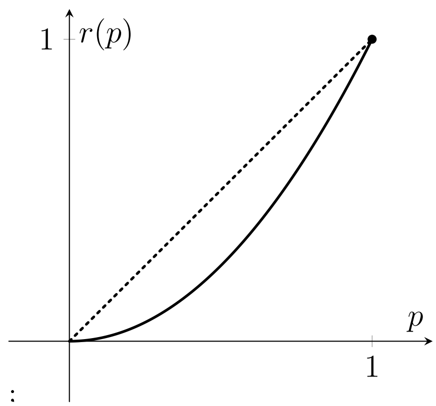

:

Figure 2: A risk function $r(p)$ that is convex with respect to probability and returns values less than p for all $0 < p < 1$.

If the agent has a risk function like this (if it is sufficiently convex), REU theory can deliver the risk-averse verdict in the above case as well. Using a linear utility function, A delivers utility 1 for sure, and B delivers utility 101 with probability one in 100. So, the REU of each option will be:

$$\text{REU}(A, r) = 1 \qquad\qquad \text{REU}(B, r) = 0 + 101r(1/100)$$

With a sufficient aversion to risk—with a sufficiently convex r function—we will obtain $r(1/100) < \frac{1}{101}$, so that B will have strictly lower REU than A. Thus, REU theory too seems able to say, for sufficiently risk-averse agents, that it is better to save one person for sure.

3 A long-run argument against EU and REU theory

In practice, each of these proposals may fail to uphold risk-averse verdicts at all. So argue various decision theorists (Thoma, 2019; Zhao, 2021, e.g.), in the form of long-run arguments. A general version of those arguments runs as follows. (This version is adapted from Thoma, 2019, but, while Thoma's argument is targeted only at REU theory, it can be applied to EU theory to similar effect.)

First, take any decision between two risky options: say, L and M, of which L's expected value is at least as great as M's. And imagine that the decision will be repeated some n times, with the value resulting from L (or M) generated independently each time.

How should we compare our options in a decision that features in such a sequence of repetitions? One rule for doing so—the most plausible rule, according to Thoma[9]—is *resolute choice*: that the best option in each decision in the sequence is whichever option appears in the sequence of choices that is best *overall*.[10] One such sequence of choices is L every time, for which the overall payoff is the sum of value of n independent trials of L. If L represented as a random variable over possible values, that sequence of choices is given by the variable $L + L + L + ... + L$ (repeated n times), or $L_{\times n}$ for short.

Which is better, the sequence of choices associated with $L_{\times n}$ or with $M_{\times n}$?[11] Expected value theory would match what it says about L and M: that the repeated version of L is at least as good as the repeated version of M; repeating the options makes no difference to the verdict. But what do risk-aversion-accommodating theories say?

There is reason to expect that any such theory will say the same as expected value theory, so long as the number n of repetitions is large enough. According to the (Strong) Law of Large Numbers[12], as the number of repetitions n approaches infinity, with probability 1, the average *actual value* obtained from each trial in $L_{\times n}$ approaches the average *expected value* of the individual options in the sequence, $\mathbb{E}(v(L))$. And the possibility that the average actual value is much less than $\mathbb{E}(v(L))$ has probability approaching 0 as n gets large. In this sense, as n becomes large enough, the option $L_{\times n}$ can be approximated more and more closely by an option that delivers value $n \times \mathbb{E}(v(L))$ for sure (and similarly for $M_{\times n}$). You might think that, therefore, $L_{\times n}$ and $M_{\times n}$ should be evaluated the same way as such options that deliver those expected values for sure. If so, we will be led to endorsing the same verdict as expected value theory, whether or not we instead endorse a theory that is thought to accommodate risk aversion.

Consider the case from earlier, of the two charities one might donate to in order to combat trachoma. (Merely for simplicity, I assume that there are no other sources of value in the world.)

A : (an outcome of) value 1 with probability 1.

B: value 101 with probability $1/100$; value 0 otherwise

Suppose that you make the same decision 10,000 times over (perhaps with the donation in B to different charities each time). What do the above risk-aversion-accommodating theories say of $A_{\times 10{,}000}$ and $B_{\times 10{,}000}$?

Take expected utility theory, and a risk-averse agent with utility function, say, $u(v) = \ln(v + 1)$ (the natural logarithm of $v+1$). When evaluating the expected utility of $A_{\times 10{,}000}$, since $A_{\times 10{,}000}$ delivers value 10,000 for sure, the agent need only convert value 10,000 into utility: approximately 9.21. And while evaluating the expected utility of $B_{\times 10{,}000}$, they have a much more complicated calculation to do, but the

[9] One key advantage of resolute choice is precisely that allows risk-averse verdicts to change when we know the decision will be repeated. If an agent faces the decision from above between A and B thousands of times over, with independent payoffs each time, they will almost certainly obtain far more value if they choose $B_{\times n}$ than $A_{\times n}$. It seems that, at least sometimes, this knowledge should justify making riskier choices rather than just choosing A every time. For related discussion, see Stefánsson (2020).

[10] This definition differs from those used elsewhere McClennan (e.g. in 1990), in that it concerns betterness rather than permissibility.

[11] As it turns out, there are no other contenders for best overall plan. This is not just a matter of intuition; it follows from Theorems 3 and 4 below. That result allows us to simplify the discussion, considering only the plans associated with $L_{\times n}$ and $M_{\times n}$ in cases like this.

[12] See Feller (1968, p. 258).

expected utility becomes 9.22—greater than for $A_{\times 10,000}$. So, for agents with this utility function, EU theory says that $B_{\times 10,000}$ is better than $A_{\times 10,000}$.

Alternatively, take REU theory, and a risk-averse agent with risk function $r(p) = p^2$, as Buchak (2013, p. 50) suggests would be reasonable (and, as throughout, a linear utility function). Their REU for $A_{\times 10,000}$ will simply be 10,000. But their REU for $B_{\times 10,000}$ becomes approximately 10,100—greater than for $A_{\times 10,000}$. In fact, the analogous result holds even when A and B are repeated just 11 times. For *any* $n \geq 11$, $B_{\times n}$ will have higher REU than $A_{\times n}$.

You might expect this to hold in general. Take *any* option L, and repeat it n times; it seems that its average EU (the EU of $L_{\times n}$, divided by n) will approach the utility of L's expected *value* as n gets large. Likewise, it seems that its average *REU* will approach its average EU. If so then, for large enough n, both EU theory and REU theory will agree with expected value theory about whether $L_{\times n}$ is at least as good as any given $M_{\times n}$. And, if we endorse resolute choice, EU theory and REU theory then imply in each decision in the sequence that L is at least as good a choice as M. These theories say just the same as expected value theory!

You might think that all of our practical decisions are taken within such sequences. Many of our decisions are repeated many times. We decide many times in our lives how to travel to work—whether to take the environmentally friendly method or the environmentally damaging one. Or you may decide many times in your life whether to donate $50 to one charity or another. (Donate a large amount at once, with the option of splitting it between two charities, and you effectively make this decision many times at once!) Or, at least, we will face other decisions with similar options, probabilities, and values. From this we might infer that REU theory and EU theory will agree with expected value theory in nearly all practical cases.

Crucially, *even if* a decision is not repeated so many times, it still seems likely that those theories will agree on the best option (Thoma, 2019, p. 248). That decision still features in a long sequence of decisions, and the best plan for that sequence will likely be to always choose the option with the highest expected value. So it seems that EU theory or REU theory (combined with resolute choice) will nearly always agree with expected value theory even in decisions that are never repeated.

Given this confluence of our best risk-aversion-accommodating theories and risk-neutral theories, why not simply accept the latter? Expected value theory is the simpler theory, after all. And deviations from it are typically justified by intuitions that it is permissible to choose in a risk-averse manner in particular cases. But, by the above argument, accommodating risk aversion within our theories in either of the above ways will not vindicate such intuitions. So, we might as well banish risk aversion. At least, so goes the argument.

4 Evaluating the argument

This argument as presented above has problems, four of which I will address below. Some of these problems can be patched over—for instance, where the argument appears to rely on speculative inferences, we can make it more rigorous. Others cannot—there remain many practical decision scenarios where EU theory and REU theory simply need not agree with expected value theory.

4.1 The law of large numbers and EU/REU

It is true that the Law of Large Numbers implies that the average value resulting from a repeated option approaches its expected value. As n approaches infinity, the probability that $L_{\times n}$'s average value is far from $\mathbb{E}v(L)$ approaches 0. So, for large enough n, the average payoff is all but guaranteed to be close to the average EU. This holds for any option L. And from this we might infer that the average EU and average REU of $L_{\times n}$ approaches the utility of $\mathbb{E}v(L)$.

But this inference is much too quick. The inferred result may hold in some specific cases, such as that described above. But why think that, in *all* cases, the average EU, average *REU*, and average expected value of a repeated lottery all behave in the same way? After all, the Law of Large Numbers says nothing about the EU or REU of an option. And, if we cannot make this inference, the argument seems to fall at the first hurdle.

Fortunately, something like the Law of Large Numbers can be proven for both EU and REU, so (something like) that inference can be defended. First, for the EU of repeated options, I prove the following. (See the appendix for all proofs.)

> **Theorem 1**: For any real-valued random variables over value, L and M, and any increasing utility function $u(v)$ with strictly increasing or strictly decreasing absolute risk aversion, there is some $n' > 0$ such that, for all greater $n > n'$, if $\mathbb{E}v(L) > \mathbb{E}v(M)$ then $\mathbb{E}u(L_{\times n}) > \mathbb{E}u(M_{\times n})$.

Take any options L and M, of which L has greater expected value. Theorem 1 states that, for any large enough n, EU and expected value theory will agree about how to compare $L_{\times n}$ and $M_{\times n}$.[13] If you faced a long enough sequence of repeated decisions between L and M, and if resolute choice holds, then the two theories would agree that L is the better option every time.

But the theorem is limited. For one, it only holds for pairs of options where one has *strictly greater* expected value. We often cannot say the same for L and M with *equal* expected value. Yes, the average EU of $L_{\times n}$ will still approach the EU of L, and likewise for M. But the verdicts of EU and expected value theory may *always* give opposite verdicts. For instance, consider the safe option A and the coin flip C:

A: value 1 with probability 1.

C: value 2 with probability $1/2$; value 0 otherwise

Both options have the same expected value: 1. But C is a riskier option and, accordingly, risk averse

[13] Note that this result is contrary to claims by Samuelson (1963) and Stefánsson (2020). Each claims something to the effect that, if EU theory disagrees with expected value theory about how to compare L and M, then they must also disagree about how to compare $L_{\times n}$ and $M_{\times n}$. But this is because they each assume that a risk-averse EU maximiser's willingness to take an additional risky bet must not depend on the value they have already secured—they assume constant absolute risk aversion, and so a linear or exponential utility function. But these represent only a small slice of the possible utility functions on offer!

agents will give it lower EU than A. What happens if we repeat A and C arbitrarily many times? It turns out that an agent with a concave utility function, such as $u(v) = \ln v$, will assign a greater EU to $A_{\times n}$ than to $C_{\times n}$, *no matter how many* times n the options are repeated. So, EU theory and expected value theory will always disagree in such cases.

For another limitation, note that Theorem 1 only holds for utility functions $u(v)$ with strictly increasing or decreasing *absolute risk aversion*. Absolute risk aversion is a technical notion from economics, describing the agent's attitude to risk at each level of value.[14] For our purposes, all that matters is that there are utility functions where this measure is neither increasing nor decreasing. For instance, linear utility functions such as $u(v) = v$ have absolute risk aversion of 0. And *exponential* utility functions—those of the form $u(v) = 1 - a^{-v}$ (for some positive a)—have constant absolute risk aversion too. With an exponential utility function, an agent would be no more nor less willing to take a bet if they had already won or lost an arbitrary number of previous bets. Because of this, it turns out that an agent with such a utility function will rank $L_{\times n}$ as better than $M_{\times n}$ only if L has greater EU than M, no matter how many times n each option is repeated.[15] So again, with a utility function such as this, EU theory and expected value theory will disagree, no matter how many times a decision is to be repeated.

I prove a similar result for the REU of repeated options.[16]

> **Theorem 2**: For any real-valued random variables L and M over value and any continuous[17] risk function r, there is some $n' > 0$ such that, for all greater $n > n'$, if $\mathbb{E}v(L) > \mathbb{E}v(M)$ then $\text{REU}(L_{\times n}, r) > \text{REU}(M_{\times n}, r)$.

As above, this theorem is limited. Again, it only holds for L and M where one has *strictly greater* expected value. When L and M have *equal* expected value, the average REU of each of $L_{\times n}$ and $L_{\times n}$ may approach the average value of L and M, respectively. But REU theory may still *never* agree with expected value theory about the ranking.

For example, consider again the case of A versus C from above. Both options have equal expected value. But a generally risk-averse agent will give C a lower REU than A. But they will also give $C_{\times n}$ a lower REU than $A_{\times n}$, no matter how many times n each option is repeated. Although the average REU of $C_{\times n}$ may approach C's expected value, with a convex r, it will never quite get there. The same goes for *any* pair of options with equal expected value and unequal REU, no matter how low their stakes. The same also goes for *any* number n of repetitions and for *any* generally risk-averse agent, no matter how mild their risk aversion. For any such agent, in any such cases, the verdict of REU theory will remain different from expected value theory.

Where does this leave long-run arguments against risk aversion? In one sense, they are strengthened—we have proof that, for at least some wide class of cases and utility functions, and sufficiently many

---

[14] It is measured by the expression $\frac{u''(v)}{u'(v)}$, where u' and u'' are, respectively, u's first and second derivatives with respect to v.

[15] This is shown by Samuelson (1963, p. 53).

[16] Both this result and Theorem 1 above are roughly equivalent to the *repetition theorem* sketched by Buchak (2013, pp. 214 & 218) and appealed to by Thoma (2019, p. 238).

[17] To streamline the discussion, I will assume from here on that REU theory requires continuous risk functions. But note that Buchak (2013, p. 49) Buchak herself does not rule out discontinuous risk function. If discontinuous risk functions are valid, then there are many more cases in which REU and EU theory can disagree.

repetitions of each case, these risk-aversion-accommodating theories do come to agree with expected value theory. But in a more important sense, they are weakened. There are mundane decision scenarios—any in which options have the same expected value but different EU or REU—in which those theories *never* come to agree with expected value theory. And there are utility functions—any with constant relative risk aversion, such as exponential utility functions—that never change their mind about how to compare options, no matter how many times they are repeated. These theories can still deliver verdicts entirely different from expected value theory.

4.2 Limits of limit results

The exceptions don't end there. Cases of equal expected value are not the only ones where EU and REU theory consistently disagree with expected value theory.

The Law of Large Numbers and the two theorems above are limit results. They may say that one quantity approaches another—that the average utility of $L_{\times n}$ approaches the utility of L's expected value, or that the average REU of $L_{\times n}$ does so—but this need only occur as n approaches infinity. By themselves, limit results like these say *nothing* about the behaviour of the average utility or the average REU for any given finite n.[18] But the argument against risk-aversion-accommodating theories requires saying something about a given finite n. After all, mortal agents like us face the same decision only finitely many times in our lives—far fewer than infinitely many times—and the argument requires that this finite number is sufficiently large. So, for mortal agents like us, those risk-aversion-accommodating theories might still give verdicts entirely different from expected value theory.

To determine whether the two theories agree in particular decisions, we must look to specific examples. For instance, consider again A and B.

> A: value 1 with probability 1.
>
> B: value 101 with probability $1/100$; value 0 otherwise

Suppose we make this comparison using REU theory. For an agent who is sufficiently risk averse, A has the higher REU. For instance, if an agent's risk attitude is such that $r(p) = p^2$ (and their utility function is linear), then $\text{REU}(B, p^2) = 0.0101$, which is a lot less than $\text{REU}(A, p^2) = 1$.

Suppose that we also use REU theory to compare repetitions of each option. It turns out to take very few repetitions for an agent with that same risk attitude to say that $B_{\times n}$ is better than $A_{\times n}$. With just 11 repetitions, the repeated risky option will have higher REU than the safe one. So, for an agents who must decide between similar options at least 11 times in their life, and whose risk attitudes are such that $r(p) = p^2$, REU theory will agree with expected value theory that it is better to choose B than A here.

But 11 repetitions will not be enough for agents with more extreme risk attitudes. For instance, take an agent so risk averse that $r(p) = p^4$. For them, the decision would need to be repeated 55 times for REU

[18] I am grateful to Alan Hájek for pointing this out.

theory to agree with expected value theory that it is better to choose B each time. And, for $r(p) = p^5$ (or p to a higher power), it would take well over 1,000 repetitions.[19]

Indeed, we can take any pair of options that can have unequal REU and expected value, and take *any* finite n, and there will be *some* risk function such that n is not high enough. So, for any such decision, there is a possible agent for whom REU theory and expected value theory will disagree.

What about EU theory? If the agent's risk attitude is such that $u(v) = \sqrt{v}$, for instance, A will have higher EU than B. And this ranking carries over to comparisons of $A_{\times n}$ and $B_{\times n}$, for n up to 2,540. Only for n higher than that do EU theory and expected value theory come to agree.

As above, for even more extreme risk attitudes, it requires even greater n for the theories to agree. For $u(v) = \ln(v+1)$, n must be at least 5,060 for $B_{\times n}$ to have greater EU than $A_{\times n}$. (And, as noted above, exponential utility functions, no finite n gives $B_{\times n}$ the greater EU.) In general, for any pair of options on which EU theory and expected value theory can disagree, and any finite n, there will be *some* utility function such that n is not high enough for them to agree. So, for any such decision, there are possible agents for whom EU theory and expected value theory will give different verdicts.

4.3 Rare decisions

The range of possible exceptions becomes even greater if consider decisions that are repeated few times or not at all.

Many of our most important decisions are like this; some arise only once. For instance, you might think that one's initial decision of which career path to pursue is sufficiently high in its stakes that they will never repeat it. Likewise, many decisions with high *moral* stakes are not decisions that most of us make more than once in our lives. For paradigmatic examples, consider Harry Truman's decision of whether or not to detonate nuclear weapons over Hiroshima and Nagasaki; or the philanthropist Katharine McCormick's decision of whether or not to fund the (eventually successful) combined oral contraceptive pill (Eig, 2015, see); or Norman Borlaug's decision of whether to research new breeds of wheat that could (and eventually did) dramatically increase the world's food supply. In these high-stakes risky decisions, perhaps risk-aversion-accommodating theories still have something to say different from expected value theory, *even if* we restrict our attention to agents with only moderate aversion to risk. And, if those theories have something different to say even just in these cases, where it is all the more important to apply the correct theory of instrumental betterness, they clearly do represent a distinct alternative to expected value theory.

The worry that some decisions go unrepeated is briefly addressed by Thoma: "we can speculate that in the case of large numbers of *different* small-stakes gambles, too, resolute REU maximizers will behave approximately like expected [value] maximizers" (Thoma, 2019, p. 248; emphasis mine). It seems plausible that it need not be the case that each exact decision is repeated. Having very many *different* decisions in one's future will likely be enough to steer our verdicts towards risk neutrality, even if many of those decisions have low stakes.

[19] As Buchak (2013, p. 214) notes, "...in general, the more risk-avoidant an agent is, the higher n will have to be for her to prefer the [Bs]".

Looking at particular sequences of lotteries, Thoma's speculation does seem correct. For instance, consider again the options A, B, and C.

A : utility 1 with probability 1.

B: utility 101 with probability 0.01; utility 0 otherwise

C: utility 2 with probability $\frac{1}{2}$; utility 0 otherwise

Suppose an agent must choose between A and B, but knows they will face n other decisions in their life in which C is unambiguously the best option. How do $A + C_{\times n}$ and $B + C_{\times n}$ compare? It turns out that, as n increases, the EU of $B + C_{\times n}$ comes to be greater than that of $A + C_{\times n}$ (at least for utility functions with strictly increasing or strictly decreasing absolute risk aversion). Likewise, the REU of $B + C_{\times n}$ comes to be greater than that of $A + C_{\times n}$ (at least for continuous risk functions). So, even if an agent faces the decision between A and B just once, both theories will still recommend B if both options would be followed by sufficiently many other decisions in which those theories recommend C. In this case at least, both theories will agree with expected value theory.

It turns out that something similar occurs for many pairs of options and for many sequences of decisions that might follow them. We can prove the following.

Theorem 3: For any real-valued random variables L and M, any ordered sequence of infinitely many lotteries $\{N_1, N_2, N_3, ...\}$ with $\text{var}(N_i) > 0$ and bounded above[20], and any increasing utility function $u : \mathbb{R} \to \mathbb{R}$ with strictly increasing or strictly decreasing absolute risk aversion, there is some $n' > 0$ such that, for all $n > n'$, if $\mathbb{E}v(L) > \mathbb{E}v(M)$ then $\mathbb{E}u(L + N_1 + N_2 + ... + N_n, r) > \mathbb{E}u(M + N_1 + N_2 + ... + N_n, r)$.

Theorem 4: For any real-valued random variables L and M, any ordered sequence of infinitely many lotteries $\{N_1, N_2, N_3, ...\}$ with $\text{var}(N_i) > 0$ and bounded above, and any continuous risk function r, there is some $n' > 0$ such that, for all $n > n'$, if $\mathbb{E}v(L) > \mathbb{E}v(M)$ then $\text{REU}(L + N_1 + N_2 + ... + N_n, r) > \text{REU}(M + N_1 + N_2 + ... + N_n, r)$.[21]

These results are similar to Theorems 1 and 2. They state that, if both options L and M are followed by a sufficiently long sequence of the same choices in other decisions, REU theory and EU theory will typically agree with expected value theory that L followed by that sequence is better than M followed by the same sequence. There are exceptions to this—where L and M have *equal* EU, where the EU-maximising agent's utility function exhibits constant absolute risk aversion, and where the REU-maximising agent's risk function is discontinuous near 0 or 1—but these are the same as for Theorems 1 and 2. So, in practice, if you ever need to decide between some options L and M with unequal expected value, and you don't

[20]This is quite a tame requirement. It will be satisfied when the sequence of options $\{N_1, N_2, N_3, ...\}$ is something like C repeated over and over. Where it fails to hold is if the agent's life is devoid of other risky decisions, or if the payoffs in those decisions are unbounded such as in the St Petersburg and Pasadena games.

[21]Theorems 3 and 4 reflect a more general phenomenon: that, as some 'background uncertainty' B increases in variance and in stakes, $L + B$ eventually comes to stochastically dominate $M + B$ (see Tarsney, 2020; Wilkinson, 2022, §6), such that any plausible decision theory will eventually favour $L + B$.

have certain problematic risk attitudes, then EU theory and REU theory will give you the same verdict as expected value theory so long as you have *sufficiently many* other decisions remaining in your life.

But how large must n be to be sufficiently many? This will depend on just how risky those later options $N_1, N_2, ..., N_n$ are, as well as on the stakes involved in the initial options L and M, and on the agent's risk attitude. Crucially, we can take any n, any $N_1, N_2, ..., N_n$, and any continuous r function, and there will be *some* options L and M such that the result does not hold. For instance, consider a decision between options D and E, where you know each will be followed by $n = 10,000$ more decisions in which you will choose C (from above) each time.

Take EU theory, and take a fairly moderate risk attitude such that $u(v) = \sqrt{v}$. Even when C is repeated 10 thousand times, $E + C_{\times 10,000}$ will still not have as great an EU as $D + C_{\times 10,000}$. Likewise, we could repeat C one million, one billion, or any number of times; there will be some ever-higher-stakes options D' and E' such that EU theory still recommends the safe option.

Or take REU theory, and the fairly moderate risk attitude such that $r(p) = p^2$. Even when C is repeated 10 thousand times, $D + C_{\times 10,000}$ will still not have as great an REU as $E + C_{\times 10,000}$. Likewise, we could repeat C one million, one billion, or any number of times; there would be some ever-higher-stakes options E' and F' such that REU theory still recommends the safe option. So we have another exception to the argument that risk-aversion-accommodating theories can say nothing different to expected value theory: decisions that are sufficiently high in their stakes. If an agent faces few *extremely* consequential decisions in their life—such as, perhaps, a political leader deciding whether to start a war, or deciding how to direct foreign aid—both EU theory and REU theory will often offer advice different from expected value theory.

So EU theory and REU theory (combined with resolute choice) agree with expected value theory in fewer cases still. Not only do they disagree in cases where options have equal expected value and cases where the agent's risk attitudes are extreme enough; they also disagree in cases where the stakes are high enough. Indeed, they disagree in high-stakes cases (like that above) *even if* the agent is only moderately risk-averse (and so does not have too convex a risk function under REU theory, or too concave a utility function under EU theory).

You might think that this further observation about high-stakes cases bears little on the case for or against these theories—that disagreeing with expected value theory in rare, high-stakes cases does little to distinguish them. (Thoma, 2019, p. 252), for instance, takes this view: that, if a "...theory ends up making approximately the same recommendations regarding *small-stakes* gambles as expected [value] theory [then it] offers no real alternative to expected [value] theory..." (emphasis mine). But this conditional does not seem plausible, especially when we apply said theory in the context of moral decision-making. In moral decisions, where the stakes may be enormous, it is especially important for our theory to give the correct verdicts in the cases with the highest stakes. Even if two theories agree in almost all practical cases, that matters little if they disagree in exactly those cases which are the most important for our theories to get right. (Imagine rejecting utilitarianism or Kantianism because they agree in most everyday decisions!) So, it seems that EU theory and REU theory *could* offer a genuine alternative to expected value theory, *even if* they only disagreed with it in high-stakes cases. And, as demonstrated above, they disagree in many more cases than that.

4.4 Abandoning resolute choice

So far I have assumed that resolute choice holds—that if $L_{\times n}$ is better than $M_{\times n}$ then, in a decision between L and M that is repeated n times, L should be considered a better option than M every time. But what if we rejected it?

We would be left with strictly less plausible rules of dynamic choice, according to both Thoma's and Zhao's versions of the argument. Each argues that there are compelling reasons to think that resolute choice is more plausible than its rivals (see Thoma 2019, pp. 242-4 and Zhao 2021, §3.1). But if you find risk aversion independently plausible, this likely won't convince you. After all, you are not deciding between resolute choice and its rivals in a vacuum. You must decide between them when accepting resolute choice means that we must deny risk-aversion-accommodating theories, or so the argument goes. If you find arguments for EU theory or REU theory more compelling than those for resolute choice, those rival dynamic choice rules are still worth considering.

But what are those rules?[22] One is *sophisticated choice*: the best option in a decision is whichever is best conditional on the agent choosing the *best* option in each of their future decisions.[23] To illustrate this rule, suppose you face a sequence of n repeated decisions between A' and B.

> A': value 1 with probability 0.9; value 0 otherwise.
>
> B: value 101 with probability 0.01; value 0 otherwise

Working backwards, in the nth decision, we can suppose that A' is the best option—that it has the greater EU or REU, according to the agent's risk attitude. Then, in the $(n-1)$th decision, you must evaluate your two options as $A' + A'$ and $B + A'$. Both options feature spreads of risk different to those of A' and B. (E.g. each has a much lower probability of 0 than they do without A' added.) Indeed, on some risk functions, the REU of $A' + A'$ ends up being *lower* than that of $B + A'$. Likewise, on some utility functions, the EU of $A' + A'$ ends up being lower too. In fact, on any utility function with increasing or decreasing absolute risk aversion, EU theory will say that $B + A'_{\times n}$ becomes better than $A' + A'_{\times n}$ once n gets large enough (by Theorem 3). Likewise, on any continuous risk function, REU theory will say that $B + A'_{\times n}$ becomes better once n gets large enough (by Theorem 4). So, both theories will say that B is the better option in the $(n-1)$th decision, as expected value theory does, if we accept sophisticated choice. And this agreement with expected value theory also holds for any options L and M (with unequal expected value) each followed by *any* sufficiently long sequence of options $N_1, N_2, ..., N_n$. So, at least for the first few decisions in any such sequence, risk-aversion-accommodating theories will agree with expected value theory, at least in some cases. The argument against those theories suddenly seems more powerful:

[22]One option I won't cover in any depth is *permissive resolute (or sophisticated) choice*: the best option in a decision is whichever would be best according to resolute (or sophisticated) choice if all of the agent's future decisions beyond some time horizon t were excluded, with t set by the agent's own characteristics. (Alternatively, we might consider a similar choice rule by which all settings of t are valid, and an option must be better than another on *all* such settings of t to be better *simpliciter*.) For compelling arguments against this option, see Thoma (2019, pp. 250-2).

[23]For an argument that we should not endorse sophisticated choice and REU theory simultaneously, see Buchak (2015, §4)).

accept resolute choice *or* sophisticated choice, and the two theories will often agree; we need not assume the former.[24]

A different choice rule you might adopt is *self-predictive choice*, which says: the best option in the decision before you is whichever is best in light of your uncertainty about what you will choose in each of those future decisions.[25] This is the rule you would use if you thought that your past choices were irrelevant to present evaluations, and that future choices relevant only insofar as your predictions of what you'll choose change your present options. To illustrate, suppose that you must decide between the original A and B options, with that decision repeated again just once. And suppose that, when facing that first decision, you are uncertain of how you will choose next time, with probability 0.5 that you will choose each option. (If you choose B, you will also be uncertain of how the option turns out; your overall uncertainty of how the second decision turns out can be denoted by L.) By self-predictive choice, you must evaluate your options in the first decision as though they are the following standalone options:

> $A + L$: value 102 with probability 0.005; value 2 with probability 0.5; value 1 with probability 0.405.
>
> $B + L$: value 202 with probability 0.00005; value 102 with probability 0.005; value 101 with probability $\sim$ 0.05; value 1 with probability 0.495; value 0 with probability $\sim$ 0.49

These options feature spreads of risk different from A and B. And, on some risk attitudes, EU theory and REU theory will judge $B + L$ as better than $A + L$, *even though* they judge A as better than B. Add sufficiently many gambles L of future payoffs and, on any utility function with increasing or decreasing absolute risk aversion and on any continuous risk function, $B + L_{\times n}$ will have greater EU and greater REU than $A + L_{\times n}$. So, EU theory and REU theory will then agree with expected value theory. And, thanks to Theorems 3 and 4 above, this generalises. The same result will hold if we replace A and B with any options with unequal expected value. Likewise, it holds if we substitute L for lots of different lotteries, even if they are far less risky (as they might be if you are almost certain of what your future self will choose).

We might adopt either sophisticated choice or self-predictive choice Either way, we can deny resolute choice and our theories can still obtain verdicts in a principled manner. And, as demonstrated here, these choice rules will often lead EU theory and REU theory to deliver the same verdicts as expected value theory. It seems that the argument against risk-aversion-accommodating theories is stronger than we might have thought.

But these dynamic choice rules still leave the exceptions described above. In addition, they introduce new exceptions. In particular, both sophisticated and self-predictive choice lead EU theory and REU theory to disagree with expected value theory in the later decisions of an agent's life. By either rule, if an agent's final decision is between L and M, they must compare them in isolation—they must judge as better whichever of those options has greater EU or REU, considered alone.[26] So either theory will

[24] Thoma (2019, p. 244) makes a similar, weaker claim: that sophisticated choice results in REU and expected value theory coming to agree in *some* cases, specifically when there is no known end to a sequence of decisions. But, in fact, the theories come to agree in a much broader class of repeated decisions, as shown here.

[25] This is the dynamic choice rule of *subjective actualism*, as advocated by Jackson (2014), among others.

[26] Note that this matches conventional investment wisdom: that it is rational to transition to a less risky portfolio as one

disagree with expected value theory about how to compare any end-of-life options that, taken alone, the risk-aversion-accommodating theory would rank differently. Likewise, by either sophisticated or self-predictive choice, we can rewind to a few decisions before the agent's final decision, and REU theory and EU theory will often disagree with expected value theory there as well—on some risk attitudes, for some small n, and/or for some high-stakes options L and M (such as, perhaps, how to allocate one's will), risk-aversion-accommodating theories will judge $L + N_1 + ... + N_n$ as better than $M + N_1 + ... + N_n$ even though expected value theory says the reverse. So, as an agent's life comes to an end, the two theories will begin to sometimes give very different verdicts, by either of these choice rules.

5 Conclusion

Where does this leave our most promising risk-aversion-accommodating theories? Do they fail to offer a true alternative to expected value theory, as long-run arguments have suggested?

Such theories do offer an alternative. As shown here, they disagree with expected value theory in various classes of cases: whenever agents compare options with *equal* expected value, and of which one is riskier; when any given pair of options are compared by agents who are simply risk-averse *enough*; when even agents who are only moderately risk-averse face decisions with particularly high stakes; and, if we reject resolute choice, then also when an agent compares options available to them near the end of their life. And these will include moral decisions on which it is especially important for our theory to get right: high-stakes decisions such as, perhaps, a political leader deciding whether to start a war, or an elderly philanthropist deciding where to bequeath their riches. It will not do to simply follow the verdicts of expected value theory in these cases if instead some risk-aversion-accommodating theory is true—disagreement even just in these cases is enough to warrant retaining such theories.

Admittedly, it is true that these theories agree with expected value theory in many cases. We now have formal proof of that. Take any decision between two options (with unequal expected value). If an agent has a *suitable* risk attitude and faces *sufficiently* many other decisions in their life, REU theory and EU theory will agree with expected value theory on the verdict.

You might think that this occasional agreement is itself reason enough to abandon those theories. Those theories are often justified by intuitions in practical cases where a particular risk-averse verdict seems plausible (see, e.g. Buchak, 2013, §1.1). For instance, suppose an agent must decide which charity to donate to, of the two charities described in the Introduction. Intuition may suggest that a typical agent, averse to risk, ought to donate to the sure thing. So we may want decision theory to accommodate this. But it appears that EU theory and REU theory cannot do so, based on the discussion above, at least not for agents with only a moderate aversion to risk, and who will donate to charity many more times in their life.

But even this weaker objection to those risk-aversion-accommodating theories does not succeed. EU theory and REU theory can still justify any such risk-averse verdict in basic motivating cases such as

approaches retirement. It turns out that such behaviour is justified *even if* an agent values money linearly.

that of the two charities, even when that decision will be followed by many others. Those theories just require more extreme risk attitudes to do so (see §4.2). And, in practice, the required risk attitudes are not so extreme at all. When we recognise that real-world agents are engaged in long sequences of repeated decisions, those are simply the attitudes required to reach mildly risk-averse verdicts in each individual decision. So, risk-aversion-accommodating theories can still affirm risk-averse verdicts in the basic cases intended to motivate them.

And, with that, the argument that these theories offer no true alternative to expected value theory (as well as even the argument that they offer only *unmotivated* alternatives to it) fails. The mere fact that a risk-averse agent's verdicts become more risk-neutral overall in the long run is not, by itself, sufficient reason to reject risk aversion altogether. Risk aversion can indeed survive the long run.

6 Appendix: Proofs

Theorem 1: For any real-valued random variables L and M, and any increasing utility function $u : \mathbb{R} \rightarrow \mathbb{R}$ with strictly increasing or strictly decreasing absolute risk aversion, there is some $n' > 0$ such that, for all greater $n > n'$, if $\mathbb{E}(L) > \mathbb{E}(M)$ then $\mathbb{E}u(L_{\times n}) > \mathbb{E}u(M_{\times n})$.

Proof: Chebysev's Inequality[27] states that, for any random variable X with defined variance var(X)

$$P\big(|X - \mathbb{E}(X)| \geq k\big) < \frac{\text{var}(X)}{k^2} \qquad \text{for all } k > 0$$

$$\text{It follows that } P\big(|L_{\times n} - n\mathbb{E}(L)| \geq n\epsilon\big) < \frac{n\text{var}(L)}{n^2\epsilon^2} \text{ for any } \epsilon > 0$$

$$\Rightarrow P\Big(\frac{|L_{\times n} - n\mathbb{E}(L)|}{n} > \epsilon\Big) < \frac{\mathbb{E}(L^2)}{n\epsilon^2}$$

Note that, for any $\mathbb{E}(L)$, $\mathbb{E}(L^2)$, and ϵ, the right hand side converges to 0 as $n \rightarrow \infty$. And for any tiny $\delta_1 > 0$ we choose, there is some $n' > 0$ such that, for any $n > n'$,

$$P\Big(\frac{|L_{\times n} - n\mathbb{E}(L)|}{n} > \epsilon\Big) < \delta_1 \qquad \text{(i)}$$

Similarly, for any tiny δ_2, there is some $n' > 0$ such that, for $n > n'$,

$$P\Big(\frac{|M_{\times n} - n\mathbb{E}(M)|}{n} > \epsilon\Big) < \delta_2 \qquad \text{(ii)}$$

For the greater of the two n's, and for any $n > n'$, both (i) and (ii) hold.

To show that if $\mathbb{E}(L) > \mathbb{E}(M)$ then $\mathbb{E}u(L_{\times n}) > \mathbb{E}u(M_{\times n})$, it suffices to show that, for some random variables L' and M', $\mathbb{E}u(L_{\times n}) \geq \mathbb{E}u(L') > \mathbb{E}u(M') \geq \mathbb{E}u(M_{\times n})$ (for all increasing $u(v)$ with strictly increasing or strictly decreasing absolute risk aversion). We can obtain such L' (and M') by taking the worst (best) possible variables that satisfy (i) (or (ii)) and that still have the same minimum (maximum) values as $L_{\times n}$ ($M_{\times n}$). Note that min(im(L)) is the lowest value that results from L with non-zero probability,

[27] See Feller (1968, p. 233).

and max(im(M)) the highest value that results from M with non-zero probability.

Let L' be a random variable such that $P(L' = x) = \begin{cases} \delta_1 \text{ for } x = n\min(\text{im}(L)) \\ 1-\delta_1 \text{ for } x = n(\mathbb{E}(L) - \epsilon) \end{cases}$

Similarly, let M' be such that $P(M' = x) = \begin{cases} 1-\delta_2 \text{ for } x = n(\mathbb{E}(M) + \epsilon) \\ \delta_2 \text{ for } x = n\max(\text{im}(M)) \end{cases}$

Since $u(v)$ must be strictly increasing, $\mathbb{E}u(L_{\times n}) \geq \mathbb{E}u(L')$ and $\mathbb{E}u(M') \geq \mathbb{E}u(M_{\times n})$. But how do $\mathbb{E}u(L')$ and $\mathbb{E}u(M')$ compare?

$$\mathbb{E}u(L') - \mathbb{E}u(M') = \delta_1 u\Big(n\min\big(\text{im}(L)\big)\Big) + (1-\delta_1)u\Big(n\big(\mathbb{E}(L)-\epsilon\big)\Big) - (1-\delta_2)u\Big(n\big(\mathbb{E}(M)+\epsilon\big)\Big) + (\delta_2)u\Big(n\max\big(\text{im}(M)\big)\Big)$$

We can choose arbitrarily small δ_1, δ_2, and ϵ, which makes this difference arbitrarily close to $u\big(n\mathbb{E}(L)\big) - u\big(n\mathbb{E}(M)\big)$. And, since $\frac{u''(v)}{u'(v)}$ is strictly increasing or strictly decreasing, $u\big(n\mathbb{E}(L)\big) - u\big(n\mathbb{E}(M)\big)$ will not be arbitrarily small. Then, since $u(v)$ is strictly increasing and $\mathbb{E}(L) > \mathbb{E}(M)$, that difference will be positive. Therefore, $\mathbb{E}u(L') > \mathbb{E}u(M')$ and, given the above, $\mathbb{E}u(L_{\times n}) > \mathbb{E}u(M_{\times n})$.□

Theorem 2: For any real-valued random variables L and M and any continuous risk function $r : [0, 1] \to [0, 1]$ (and any increasing utility function $u : \mathbb{R} \to \mathbb{R}$ with strictly increasing or strictly decreasing absolute risk aversion), there is some $n' > 0$ such that, for all $n > n'$, if $\mathbb{E}(L) > \mathbb{E}(M)$ then $\text{REU}(L_{\times n}, r) > \text{REU}(M_{\times n}, r)$.

Proof: As above, Chebysev's Inequality implies that, for any $\epsilon > 0$,

$$\Rightarrow P\Big(\frac{|L_{\times n} - n\mathbb{E}(L)|}{n} > \epsilon\Big) < \frac{\mathbb{E}(L^2)}{n\epsilon^2}$$

Again, the right hand side converges to 0 as $n \to \infty$. And for any tiny $\delta_1, \delta_2 > 0$ we choose, there is some $n' > 0$ such that, for any $n > n'$, both (i) and (ii) from above hold.

Define L' and M' as above. Since r must be non-decreasing and u must be increasing, for any $r(p)$ and $u(v)$, $\text{REU}(L_{\times n}, r) \geq \text{REU}(L', r)$ and $\text{REU}(M', r) \geq \text{REU}(M_{\times n}, r)$. But how do $\text{REU}(L', r)$ and $\text{REU}(M', r)$ compare?

$$\begin{aligned}\text{REU}(L', r) - \text{REU}(M', r) = {} & u\Big(n\min\big(\text{im}(L)\big)\Big) + r(1-\delta_1)u\Big(n\Big(\mathbb{E}(L) - \epsilon - \min\big(\text{im}(L)\big)\Big)\Big) \\ & -u\Big(n\big(\mathbb{E}(M)+\epsilon\big) - r(\delta_2)u\Big(n\Big(\max\big(\text{im}(M)\big) - \mathbb{E}(M) - \epsilon\Big)\end{aligned}$$

Again, we can choose arbitrarily small δ_1, δ_2, and ϵ, which makes the above difference arbitrarily close to $u\big(n\mathbb{E}(L)\big) - u\big(n\mathbb{E}(M)\big)$, which is positive since $\mathbb{E}(L) > \mathbb{E}(M)$, $u(v)$ is strictly increasing, and $\frac{u''(v)}{u'(v)}$ is strictly increasing or strictly decreasing. So, $\text{REU}(L', r) >\text{REU}(M', r)$. Given the above, this implies that $\text{REU}(L_{\times n}) >\text{REU}(M_{\times n})$ for sufficiently large n.□

Theorem 3: For any real-valued random variables L and M, any ordered sequence of infinitely many lotteries $\{N_1, N_2, N_3, ...\}$ with $\text{var}(N_i) > 0$ and bounded above, and any increasing utility function $u : \mathbb{R} \to$

$\mathbb{R}$, there is some $n' > 0$ such that, for all $n > n'$, if $\mathbb{E}(L) > \mathbb{E}(M)$ then $\mathbb{E}u(L + N_1 + N_2 + ... + N_n, r) > \mathbb{E}u(M + N_1 + N_2 + ... + N_n, r)$.

Proof: The Kolmogorov Criterion[28] states that, for an ordered sequence of infinitely many random variables $\{X_1, X_2, ...\}$, if the sum $\sum_{i=1}^{\infty} \frac{\text{var}(X_i)}{i^2}$ converges, then the Strong Law of Large Numbers applies to that sequence. So, for any $\epsilon, \delta > 0$, there is some $n' > 0$ such that, for all $n > n'$,

$$P\Big(\frac{|X_1 + X_2 + ... + X_n - \sum_{i=1}^{n} \mathbb{E}(X_i)|}{n} > \epsilon\Big) < \delta$$

Consider the ordered sequence $\{L, N_1, N_2, ..., N_n, ...\}$. Since $\text{var}(N_i)$ is bounded above, $\frac{\text{var}(L)}{1} + \sum_{i=1}^{\infty} \frac{\text{var}(N_i)}{(i+1)^2}$ converges to some finite sum. Kolmogorov's Criterion then implies that, for any $\epsilon, \delta_1 > 0$, there is some $n' > 0$ such that, for all $n > n'$,

$$P\Big(\frac{|L + N_1 + ... + N_n - \mathbb{E}(L) - \sum_{i=1}^{n} \mathbb{E}(N_i)|}{n} > \epsilon\Big) < \delta_1 \qquad \text{(iii)}$$

Similarly, for any $\epsilon, \delta_1 > 0$, there is some $n' > 0$ such that the analogous result holds for $M + N_1 + N_2 + ... + N_n$ for all $n > n'$.

To show that $\mathbb{E}(L) > \mathbb{E}(M)$ implies that $\mathbb{E}u(L+N_1+N_2+...+N_n, r) > \mathbb{E}u(M+N_1+N_2+...+N_n, r)$, the method is similar to the proof of Theorem 1.

Let L'' be a random variable such that $P(L'' = x) = \begin{cases} \delta_1 \text{ for } x = \min\big(\text{im}(L)\big) + \sum_{i=1}^{n} \min\big(\text{im}(N_i)\big) \\ 1 - \delta_1 \text{ for } x = \mathbb{E}(L) + \sum_{i=1}^{n} \mathbb{E}(N_i) - \epsilon \end{cases}$

Similarly, let M'' be a random variable such that $P(M'' = x) = \begin{cases} 1 - \delta_2 \text{ for } x = \mathbb{E}(M) + \sum_{i=1}^{n} \mathbb{E}(N_i) + \epsilon \\ \delta_2 \text{ for } x = \max\big(\text{im}(M)\big) + \sum_{i=1}^{n} \max\big(\text{im}(N_i)\big) \end{cases}$

Since every utility function $u(v)$ is non-decreasing, $\mathbb{E}u(L+N_1+...+N_n, r) \geq \mathbb{E}u(L'', r)$ and $\mathbb{E}u(M'', r) \geq \mathbb{E}u(M + N_1 + ... + N_n, r)$. And by the same procedure as for Theorem 1, it can easily be shown that $\mathbb{E}u(L'', r) > \mathbb{E}u(M'', r)$ for all n' greater than some n, as required. □

Theorem 4: For any real-valued random variables L and M, any ordered sequence of infinitely many lotteries $\{N_1, N_2, N_3, ...\}$ with $\text{var}(N_i) > 0$ and bounded above, and any continuous risk function $r : [0, 1] \to [0, 1]$ (and any increasing utility function $u : \mathbb{R} \to \mathbb{R}$ with strictly increasing or strictly decreasing absolute risk aversion), there is some $n' > 0$ such that, for all $n > n'$, if $\mathbb{E}(L) > \mathbb{E}(M)$ then $\text{REU}(L + N_1 + N_2 + ... + N_n, r) > \text{REU}(M + N_1 + N_2 + ... + N_n, r)$.

Proof: As above, the Kolmogorov Criterion implies that, for any $\epsilon, \delta_1 > 0$, there is some $n' > 0$ such that, for all $n > n'$,

$$P\Big(\frac{|L + N_1 + ... + N_n - \mathbb{E}(L) - \sum_{i=1}^{n} \mathbb{E}(N_i)|}{n} > \epsilon\Big) < \delta_1 \qquad \text{(iii)}$$

Similarly, for any $\epsilon, \delta_1 > 0$, there is some $n' > 0$ such that the analogous result holds for $M + N_1 +$

[28] See Feller (1968, p. 259).

$N_2 + ... + N_n$ for all $n > n'$.

Define L'' and M'' as above. Since every risk function r is non-decreasing and every utility function $u(v)$ is strictly increasing, $\text{REU}(L+N_1+...+N_n, r) \geq \text{REU}(L'', r)$ and $\text{REU}(M'', r) \geq \text{REU}(M+N_1+...+N_n, r)$. And similarly to above, $\text{REU}(L'', r) > \text{REU}(M'', r)$ for all n' greater than some n, as required. □

References

BARON, J., 2000. *Thinking and Deciding*. Cambridge University Press, Cambridge. (cited on page 2)

BRIGGS, R., 2019. Normative theories of rational choice: Expected utility. In *The Stanford Encyclopedia of Philosophy* (Ed. E. N. ZALTA). (cited on page 2)

BUCHAK, L., 2013. *Risk and Rationality*. Oxford University Press, Oxford. (cited on pages 2, 3, 6, 8, 10, and 15)

BUCHAK, L., 2015. Revisiting risk and rationality: A reply to pettigrew and briggs. *Canadian Journal of Philosophy*, 45, 5-6 (2015). (cited on page 13)

EIG, J., 2015. *The Birth of the Pill: How Four Crusaders Reinvented Sex and Launched a Revolution*. W. W. Norton & Company, New York. (cited on page 10)

FELLER, W., 1968. *An Introduction to Probability Theory and Its Applications (3rd ed.)*. Wiley, New York. (cited on pages 5, 16, and 18)

HÁJEK, A., 2021. Risky business. *Philospohical Perspectives*, 35, 1 (2021), pp. 189–205. (cited on page 2)

JACKSON, F., 2014. Procrastinate revisited. *Pacific Philosophical Quarterly*, 95 (2014), pp. 634–47. (cited on page 14)

KNIGHT, F., 1921. *Risk, Uncertainty and Profit*. Houghton Mifflin Company, Boston. (cited on page 1)

MCCLENNAN, E., 1990. *Rationality and Dynamic Choice*. Cambridge University Press, Cambridge. (cited on page 5)

SAMUELSON, P., 1963. Risk and uncertainty: A fallacy of large numbers. *Scientia*, 98 (1963). (cited on pages 7 and 8)

STEFÁNSSON, H. O., 2020. The tragedy of the risk averse. *Erkenntnis*, (2020). (cited on pages 2, 5, and 7)

TARSNEY, C., 2020. Exceeding expectations: Stochastic dominance as a general decision theory. *Global Priorities Institute Working Paper Series*, 3 (2020). (cited on page 11)

THOMA, J., 2019. Risk aversion and the long run. *Ethics*, 129, 2 (2019), p. 230–53. (cited on pages 2, 4, 6, 8, 10, 12, 13, and 14)

WILKINSON, H., 2022. In defence of fanaticism. *Ethics*, 132, 2 (2022), p. 445–77. (cited on page 11)

ZHAO, M., 2021. Ignore risk; maximize expected moral value. *Noûs*, (2021). (cited on pages 2, 4, and 13)