Risk-Taking and Tie-Breaking

Ryan Doody
Brown University

Forthcoming in *Philosophical Studies*. Please cite the published version.

Abstract

When you are indifferent between two options, it's rationally permissible to take either. One way to decide between two such options is to flip a fair coin, taking one option if it lands heads and the other if it lands tails. Is it rationally permissible to employ such a tie-breaking procedure? Intuitively, yes. However, if you are genuinely risk-averse—in particular, if you adhere to Risk-Weighted Expected Utility Theory (Buchak, 2013) and have a strictly convex risk-function—the answer will often be no: the REU of deciding by coin-flip will be lower than the REU of choosing one of the options outright (so long as at least one of the options is a nondegenerate gamble). This turns out to be a significant worry for Risk-Weighted Expected Utility Theory. I argue that it adds real bite to established worries about diachronic consistency afflicting views, like Risk-Weighted Expected Utility Theory, that violate Independence. And that, while these worries might be surmountable, surmounting them comes at a price.

1 Introduction

Instrumental rationality is about taking the best means to one's ends. According to the orthodox position—expected utility theory (EUT)—taking the best means to one's ends involves two sorts of evaluations: first, there are facts about what one's ends are, and the extent to which one values them (information that's encoded in an agent's utility function); second, there are the facts about how effective one's means might be at realizing one's ends (information that's encoded in an agent's credence function). The value of an option is its expected utility: roughly, the weighted average of how good or bad its potential outcome might be, where the weights correspond to the agent's credences in those outcomes resulting from its performance. According to the orthodoxy, rational agents maximize expected utility.[1]

Some philosophers—notably, Lara Buchak (e.g. Buchak, 2013)—argue that EUT is unduly restrictive: it doesn't appropriately take into account the agent's *attitude toward risk*. EUT, effectively, treats agents as if they were *risk-neutral* in virtue of the fact that it only takes *local* features of a gamble into account. But, argues Buchak, it's rationally permissible for agents to care about a gamble's *global features*—the way its potential outcomes are situated in the space of possibilities (e.g., its minimum, its maximum, its variance, its spread, etc.)—in a way that EUT cannot properly accommodate.

In response, Buchak defends *risk-weighted expected utility theory* (REUT), which generalizes EUT by adding a third parameter into the evaluation of a gamble's subjective value. In addition to a utility function (representing how the agent values her ends) and a probability function (representing what the agent thinks might happen), REUT represents an agent's attitude toward risk by attributing to her a *risk function*, r. Unlike EUT, which weights the value of each potential outcome by that outcome's probability, REUT weights the value of each potential outcome by a function r of the probability of getting something at least as valuable. If r is convex, the value of outcomes above the minimum will contribute less to the overall value of the gamble. Agents with convex risk functions are, using Buchak's terminology, *risk-avoidant*: the value they will assign to a risky gamble will be lower than its expected value.

Although Buchak may well be right that EUT is inadequate, I will argue that REUT has some rather unpalatable consequences of its own. First, it's not obvious that it correctly represents what it is to be risk-averse: REUT will, I argue, sometimes undervalue gambles that it shouldn't. Second, and relatedly, REUT has some counterintuitive consequences regarding tie-breaking. When you're indifferent between two options, it's rationally permissible to take either. One way to decide between two indifferent options is to flip a fair coin, taking the one if it lands heads and the other if it lands tails. Offhand, it seems rationally permissible to employ such a tie-breaking procedure. However, if you are risk-avoidant,

[1] Presented this way, EUT sounds like a view about how one ought to rank their options given facts about one's utilities and credences. And some might balk at characterizing the view—at least presented in this way—as orthodoxy given that many decision theorists understand EUT to be primarily a view about what it is for one's ordinal *preferences* (over both outcomes and gambles) to be *consistent* and thus understand facts about one's "utilities" to be derivable from—or perhaps just shorthand for—facts about one's ordinal preferences over gambles. There are various positions regarding the relationship between one's preferences and one's utilities (see Dreier (1996) for a helpful discussion), but nothing I argue for here turns on which is correct. I've presented EUT as I have in order to highlight the ways in which it differs from REUT. Both views can be reinterpreted as views about what consistency constraints to place on one's ordinal preferences, but doing so complicates the presentation in unhelpful ways. Thanks to an anonymous reviewer for pressing this point.

there will be cases in which the value of deciding by coin-flip will be lower than the value of choosing one of the options outright. And so, in such cases, REUT says that it is *not* permissible to employ a tie-breaking procedure.

In the next section, I will present REUT using a couple of examples. In section 3, I will argue that these examples present a challenge for Buchak's claim that REUT faithfully captures what it is to be risk-averse. In section 4, I will argue that REUT offers counterintuitive guidance in cases of tie-breaking.

2 Risk-weighted Expected Utility Theory

The best way to understand REUT is to contrast it with EUT. Let's look at what EUT says in more detail.

First, some terminology. I'll use italicized lowercase letters (like *f*, *g*, and *h*, etc.) to refer to *gambles*, which are functions from events to outcomes.[2] Events (denoted with subscripted capitalized '*E*'s) are descriptions of how the world might be, which are detailed enough to capture everything the agent might care about. They are understood to be mutually exclusive and jointly exhaustive. They are the objects of an agent's *credences*, which is a probability function (written as '*Cr*') representing the agent's subjective uncertainty about how the world is. Outcomes (denoted by subscripted lowercase '*x*'s) are descriptions of all the things the agent cares about in a particular situation that could obtain. They are the inputs of an agent's utility function (written as '*u*'), which represent the agent's non-instrumental preferences. I will, for the sake of readability, use 'value' and 'utility' interchangeably.

Let $h = \{x_1, E_1; x_2, E_2; \ldots x_n, E_n\}$ be a gamble that yields, for each $1 \leq i \leq n$, an outcome x_i if event E_i obtains, and is such that $u(x_1) \leq u(x_2) \leq \cdots \leq u(x_n)$.

Expected Utility

$$\begin{aligned} EU(h) &= \sum_{i=1}^{n} Cr(E_i) \cdot u(x_i) \\ &= u(x_1) + \left(\sum_{i=2}^{n} Cr(E_i) \right) \cdot (u(x_2) - u(x_1)) + \cdots + Cr(E_n) \cdot (u(x_n) - u(x_{n-1})) \end{aligned}$$

EUT says that you ought to maximize expected utility. The expected utility of a gamble is the probability-weighted average of the values of its potential outcomes.

[2] Following Buchak (2013), I present the views using Savage's framework (Savage, 1954). I don't intend for my objections to depend on anything peculiar to this framework, however, and I assume it only for presentational convenience.

Here's one way to calculate the expected utility of a gamble. First, weight the value of each of its potential outcomes by the probability of that outcome occurring. Then, take the sum of those probability-weighted values.

Here's a different, equally as accurate, way of calculating a gamble's expected utility. First, order the gamble's potential outcomes from worst (least preferred) to best (most preferred). The expected utility of the gamble is *at least* as high as the value of its worst potential outcome. So, start with the value of the worst potential outcome. Then add to this minimum value the difference between it and the next highest potential value (i.e, the second-worst value), weighted by the probability of getting at least that amount. Then add to *that* the difference between the second-worst value and the next highest potential value (i.e., the third-worst value), weighted by the probability of getting something at least as valuable as it. And so on and so forth, until we reach the best potential outcome.

When there are only two outcomes, it's easy to see that these two methods coincide. Let x_1 and x_2 be the worst and best outcomes, respectively. And let p be the probability of x_2 occurring. The expected utility of such a gamble is: $p \cdot u(x_2) + (1 - p) \cdot u(x_1)$. This expression can be rewritten as: $u(x_1) + p \cdot (u(x_2) - u(x_1))$, which is the minimum value of the gamble plus the amount you might gain above that minimum weighted by the probability of realizing that gain.

The *risk-weighted* expected utility of a gamble can be calculated in an analogous way—with one crucial difference: instead of weighting the potential gains by their probabilities, REUT weights these potential gains by a *function* of their probabilities. So, to use the example from the previous paragraph, the risk-weighted expected utility of the gamble would be: $u(x_1) + r(p) \cdot (u(x_2) - u(x_1))$. If r is convex, $r(p) < p$ for all non-extremal values. And so, the amount you might gain above the minimum will contribute less to the overall instrumental value of the gamble than it does on EUT.

Here's what the view says in general. Again, let $h = \{x_1, E_1; x_2, E_2; \ldots x_n, E_n\}$ be a gamble that yields, for each $1 \leq i \leq n$, an outcome x_i if event E_i obtains, and is such that $u(x_1) \leq u(x_2) \leq \cdots \leq u(x_n)$.[3]

[3] It might be worried that, because REUT deviates from the traditional view by jettisoning the Independence Axiom (or, in Savage's framework, the Sure-Thing Principle), the agent will fail to have a cardinal utility function, u, that's well-defined. Independence (and the like) is what ensures that the value an outcome contributes to the value of a gamble is *separable*: it doesn't depend on which *other* outcomes might result from from the gamble. Separability plays an important role in the representation of an agent's preferences with a cardinal utility function because, very roughly, it allows us to identify an outcome's utility with how much the agent would risk to secure it. And so it might be worried that without separability we aren't warranted in representing the agent's non-instrumental preferences with a cardinal utility function.

This is a legitimate worry, but one that can be assuaged. Buchak (2013, ch. 3) (drawing on the work of Köbberling and Wakker, 2003; Machina and Schmeidler, 1992) proves a representation theorem for REUT, which shows that, if one's preferences satisfy a number of constraints (ones collectively weaker than those underlying the representation theorems for EUT), there will be a

Risk-Weighted Expected Utility

$$REU(h) = u(x_1) + r\left(\sum_{i=2}^{n} Cr(E_i)\right) \cdot (u(x_2) - u(x_1)) + \dots + r\left(Cr(E_n)\right) \cdot (u(x_n) - u(x_{n-1}))$$
$$= u(x_1) + \sum_{j=2}^{n}\left(r\left(\sum_{i=j}^{n} Cr(E_i)\right) \cdot \left(u(x_j) - u(x_{j-1})\right)\right)$$

REUT is a *generalization* of EUT: the two views coincide when $r(p) = p$, for all probabilities p. The risk function is subject to the following constraints: for all p, $0 \leq r(p) \leq 1$; $r(0) = 0$ and $r(1) = 1$; r is non-decreasing. For the sake of concreteness, let's look at a specific convex risk function: $r(p) = p^2$. (This is Buchak's go-to example of a risk function characterizing risk-aversion.)

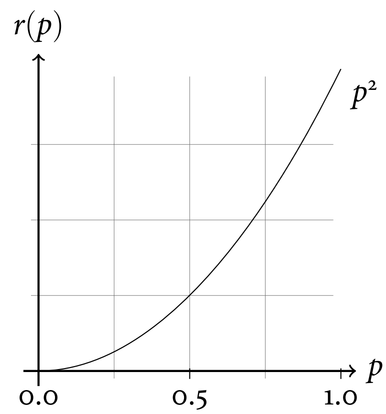

The risk function "measures how an agent structures the potential realization of some of his aims," (Buchak, 2013, p. 54). In order to better see how agents with the risk function $r(p) = p^2$ structure the potential realization of their aims, let's look at a couple of examples.

You're at the racetrack placing bets on the horses. You are considering whether to bet *for* Easy Street (f), which pays out 4 utils if Easy Street wins (E) and 2 utils if she doesn't; or to bet '*gainst* Easy Street (g), which pays out 3 utils if Easy Street doesn't win ($\neg E$) and 1 util if she does. Let's suppose that your credence in E is .25, your credence in $\neg E$ is .75, and your risk function is $r(p) = p^2$.

unique credence function, a unique risk function, and a unique (up to positive affine transformations) utility function such that one can be represented as maximizing risk-weighted expected utility relative to those functions. The result illustrates that full separability isn't necessary for an agent's utility function to be well-defined; instead, REUT only requires separability to hold among comonotonic gambles (i.e., ones that rank all events the same way), so that the contribution that an outcome x_i makes to the value of a gamble doesn't depend on which other outcomes might result *unless* those outcomes affect the relative ranking x_i occupies in the gamble. (Thanks to an anonymous reviewer for raising this worry.)

	E (1/4)	$\neg E$ (3/4)
f	4	2
g	1	3

First, note that both gambles have the same expected utility: 2½.

$$\begin{aligned} EU(f) &= 2 + \frac{1}{4} \cdot (4 - 2) \\ &= 2 + \frac{1}{4} \cdot 2 = 2\tfrac{1}{2} \end{aligned}$$

$$\begin{aligned} EU(g) &= 1 + \frac{3}{4} \cdot (3 - 1) \\ &= 1 + \frac{3}{4} \cdot 2 = 2\tfrac{1}{2} \end{aligned}$$

Interestingly, when $r(p) = p^2$, both gambles also have the same *risk-weighted* expected utility: 2⅛. Because $r(p) = p^2 \leq p$ (representing someone who is risk-avoidant), as we should expect, these gambles have lower REU than EU.

$$\begin{aligned} REU(f) &= 2 + \left(\frac{1}{4}\right)^2 \cdot (4 - 2) \\ &= 2 + \frac{1}{16} \cdot 2 = 2\tfrac{1}{8} = 2.125 \end{aligned}$$

$$\begin{aligned} REU(g) &= 1 + \left(\frac{3}{4}\right)^2 \cdot (3 - 1) \\ &= 1 + \frac{9}{16} \cdot 2 = 2\tfrac{1}{8} = 2.125 \end{aligned}$$

If you take gamble *f*, you are guaranteed to get something at least as valuable as 2 utils, and you have a 25% chance of getting something 2 utils of value more valuable. If you're risk-avoidant, the potential improvements above the guaranteed minimum are "discounted", contributing less to the overall instrumental value of the gamble than it does relative to EUT (see Figure 1). For EUT, the 25% chance of improving above the minimum adds .5 units of value; for REUT, it only adds .125 units of value. (In Figure 1, each square represents .125 utils of value. The lighter gray represents the gamble's EU and the darker gray represents its REU. For EUT, the chance of improvement above the minimum is worth four .125-util-squares, which totals .5 utils. For REUT, however, it's only worth two half .125-util-squares, which totals .125 utils.)

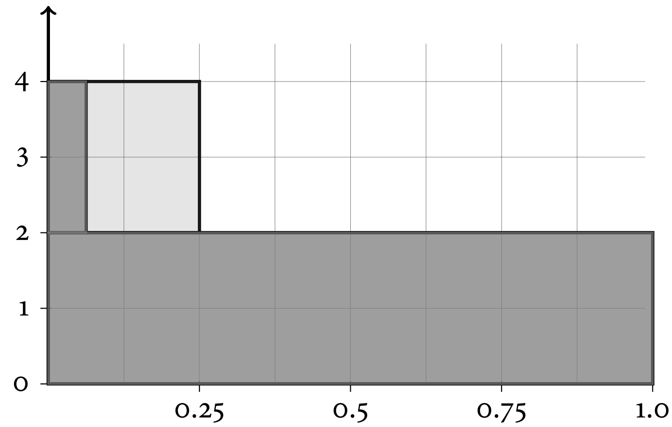

Figure 1: The risk-weighted expected utility of gamble *f* (2.125).

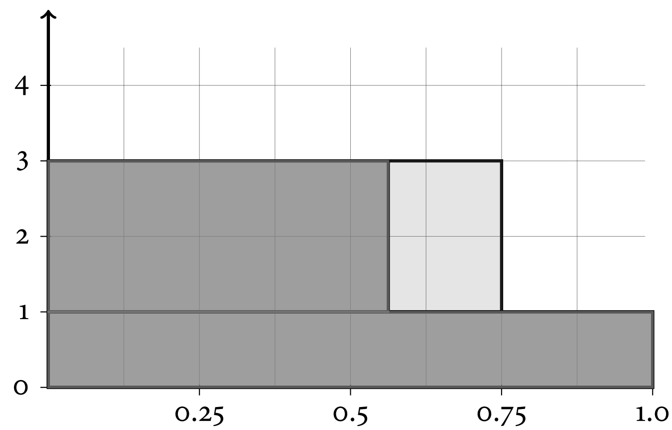

Figure 2: The risk-weighted expected utility of gamble *g* (2.125).

If you take gamble *g*, however, you are only guaranteed to get something at least as valuable as 1 util, but you have a larger chance—75% rather than 25%—of getting something 2 utils of value more valuable than that minimum. Again, if you're risk-avoidant, this potential improvement contributes less to the overall instrumental value of the gamble than it does relative to EUT (see Figure 2). For EUT, the 75% chance of improving above the minimum adds 1.5 units of value; for REUT, it only adds 1.125 units of value. (As Figure 2 shows, for EUT, the chance of improvement above the minimum is worth twelve .125-util-squares, totaling 1.5 utils. And, for REUT, that same chance of improvement is only worth eight full .125-util-squares and two half .125-util squares, totaling 1.125 utils.)

Although gamble *f* has a higher minimum than gamble *g*, gamble *g* affords a greater chance of improvement beyond that minimum. Somewhat surprisingly, according to *both* EUT and REUT, these two features balance out, making the two gambles equally valuable. Both views agree that you should be indifferent between betting for and betting 'gainst Easy Street. While the two views agree that gamble g and gamble f have the same value, they disagree about what that value is. EUT values the two gambles at 2.5, while REUT values them at 2.125.

Suppose that Eunice is an expected utility maximizer and her sister, Reu, is a risk-weighted expected utility maximizer (with the risk function $r(p) = p^2$). Suppose they share the same values and beliefs. (Also, for the sake of notational simplicity, let's suppose that they value money linearly: $u(\$x) = x$.) Eunice will be indifferent between gamble f, gamble g, and \$2.50. Reu, on the other hand, will be indifferent between gamble f, gamble g, and \$2.125. Let's say that if S is indifferent between some gamble and $\$x$, $\$x$ is that gamble's *sure-thing cash equivalent* for S. \$2.50 is gamble f's and g's sure-thing cash equivalent for Eunice; \$2.125 is gamble f's and g's sure-thing cash equivalent for Reu. Eunice would be willing to pay more to buy one of the two gambles than Reu would.

Suppose, instead of paying the bookie for one of the gambles, Eunice and Reu are offered the opportunity to select one for free. What, respectively, should they do? They are both indifferent between the two, so it's permissible for them to choose either. In the face of indifference, the sisters typically decide by flipping a fair-coin: if the coin lands heads, take gamble f; if it lands tails, take gamble g. Call this a *tie-breaking procedure*. Is it rationally permissible to employ a tie-breaking procedure in this case? It depends on the sister. For Eunice, it is. Doing so has the same expected utility as selecting either of the two gambles outright. For Reu, on the other hand, it is *not*. The 50/50 "mixture" of the two gambles has lower risk-weighted expected utility than selecting one of the gambles outright.

In the next section, we'll take a closer look at why the 50/50 "mixture" of the two gambles—let's call it $f \oplus_{1/2} g$—has lower risk-weighted expected utility than either of the two gambles themselves. I'll argue that this is a counterintuitive consequence of REUT; one that raises questions about the view's success in accurately characterizing what it is to be risk-averse.

3 Mean, Variance, and Risk-aversion

Reu, like Eunice, is indifferent between gamble f and gamble g. Unlike Eunice, it's not permissible for Reu to employ a 50/50 tie-breaking procedure—like, using the flip of a fair-coin—to make her selection.

Decide by Coin Flip

	Heads		Tails	
	E	$\neg E$	E	$\neg E$
$f \oplus_{1/2} g$	4	2	1	3

Employing such a procedure, which corresponds to a 50/50 "mixture" between the two gambles, has lower risk-weighted expected utility than either of the gambles themselves. This can be demonstrated by doing the calculation:

$$\begin{aligned} REU\left(f \oplus_{1/2} g\right) &= 1 + \left(\frac{7}{8}\right)^2 \cdot (2-1) + \left(\frac{4}{8}\right)^2 \cdot (3-2) + \left(\frac{1}{8}\right)^2 \cdot (4-3) \\ &= 1 + \frac{49}{64} + \frac{16}{64} + \frac{1}{64} = 2^1/_{32} = 2.03125 \end{aligned}$$

This might seem, however, to be somewhat counterintuitive. Because f and g have the same expected value (or, *mean*), their probabilistic mixture does as well. If their mixture is less valuable, then, it might seem like this must be because it is more risky than either of the gambles considered alone. But how might "mixing" together the two gambles introduce more risk? If anything, it seems like mixing two gambles together should create something *less* risky![4]

It might be helpful, then, in the service of evaluating REUT, to say more about what risk *is* and how it can be measured. One common way of measuring risk is with *variance*. In general, the variance of a distribution measures the extent to which its values deviate from the average. The variance of a gamble, in particular, measures how far away from its expected value the value of its potential outcomes are. There is a sense, then, in which the higher the variance, the riskier the gamble.[5]

---

[4] Consider, for example, the oft-touted financial advice to *diversify* one's portfolio. The wisdom behind the advice is that diversifying—i.e., spreading one's money among various different investments—reduces risk. And, indeed, because equally dividing a fixed amount of wealth between independently and identically distributed investments results in a portfolio with the same expected monetary return but a lower variance than those of the individual investments it contains, in many cases, diversifying reduces risk (Markowitz, 1952, 1959). (See Samuelson (1967) for a different and more general account of how, and under what circumstances, diversification is beneficial to risk-averse agents.) Because investments (in virtue of paying out different sums of money in different scenarios) are gambles, diversification might seem akin to taking their "mixture". One important difference, however, is that in classic cases of diversification, by purchasing some percentage of each investment, you might benefit from positive returns from all. The probabilistic mixture of some gambles, on the other hand, affords you no such possibility—unlike a diverse portfolio, you aren't taking some percentage of each gamble, but instead making it such that there's some chance of getting the entirety of one.

[5] Nearly all textbooks in finance, risk management, and the like include discussions of variance (and its close cousin, standard deviation) as a measure of risk. This is undoubtedly due in part to the popularity of Modern Portfolio Theory (Markowitz, 1952, 1959), which employs a "mean-variance framework" for evaluating investment portfolios. On this approach, an investment is evaluated along two dimensions: its expected return and its degree of risk (which is captured by its variance). Given this framework, Markowitz derives the set of optimal investment portfolios for different levels of risk.

Markowitz's approach has been robustly criticized. Some of the criticism is to its use of *variance* as the measure of a gamble's riskiness; some to the assumption that the value of a gamble can be neatly factored into its mean and its riskiness—however it's measured (see, for example,

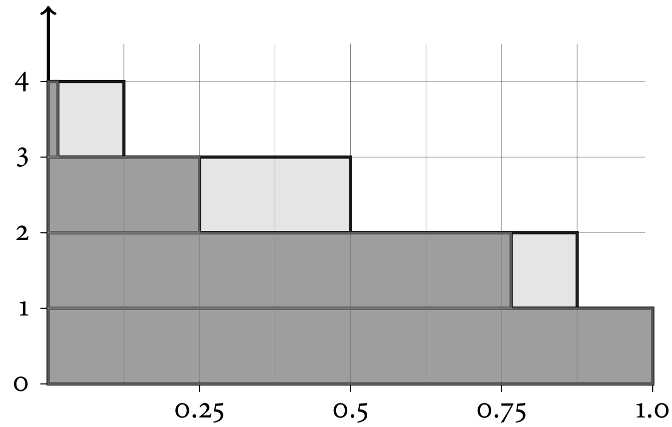

Figure 3: The risk-weighted expected utility of a fifty-fifty lottery between gamble *f* and gamble *g* (2.03125).

Let $h = \{x_1, E_1; x_2, E_2; \dots x_n, E_n\}$ be a gamble, and let $EU(h)$ be its expected value. Then, the variance of *h* is:

$$\textsc{Var}(h) = \sum_{i=1}^{n} Cr(E_i) \cdot (u(x_i) - EU(h))^2$$

Roughly, the variance of a gamble is the probabilistically-weighted average of the "distance" between each of its potential outcomes and its expected value.

Is $f \oplus_{1/2} g$ riskier—in the sense of having a higher *variance*—than gamble *f* and gamble *g*? The answer is: no. In addition to having the same mean, the three also have the same variance: 0.75.[6] If the risk of a gamble is captured by its variance, the three are all equally risky. What, then, justifies Reu in assigning a lower instrumental value to the one than to the other two?

A gamble's variance, however, is not the only way of measuring its riskiness. And, indeed, there are good reasons to worry that it is, at best, an imperfect measure of risk.[7] Instead, we could compare the riskiness of two gambles by checking to see if the one is a *mean-preserving spread* of the other. In fact, this is how Buchak defines what it is to be risk-averse. She says,

> [A]n agent is *generally risk-averse in money* (or any quantitative good)

Oddie and Milne, 1991, p. 58-9). In particular, Borch (1969) argued that, because two-dimensional mean-variance indifference curves do not exist (they cannot represent the preferences of a rational agent), the mean-variance framework is logically incoherent. For a fascinating critical discussion of these issues, see Johnstone and Lindley (2013).

[6] For the calculations and further discussion, see appendix A.

[7] See, for example, Rothschild and Stiglitz (1970, p. 241-2), and further discussion in appendix A.

> just in case of any two gambles f and g such that g can be obtained from f by a mean-preserving spread, the agent weakly prefers f; that is, of any two gambles with the same average monetary value, the agent weakly prefers the one that is less spread out, if such a judgment can be made, (Buchak, 2013, p. 21-2).

Roughly, B is a mean-preserving spread of A if, they have the same mean, and B can be obtained from A by adding "noise". Impressionistically, mean-preserving spreads shift value away from the center of a distribution out toward its tails.[8] If one gamble is a mean-preserving spread of another, it's riskier.

However, the relation "is a mean-preserving spread of" induces only a partial ordering on gambles. Obviously, if two gambles have different means, neither can be a mean-preserving spread of the other. Yet, intuitively, one might be more risky than the other. But even among gambles with the same mean, it might be that neither is a mean-preserving spread of the other. Our three gambles—f, g, and $f \oplus_{1/2} g$—are a case in point. If B is a mean-preserving spread of A, B will have a higher variance than A. But our three gambles all have the same mean and the same variance, so none are mean-preserving spreads of the other. The lower instrumental value that Reu assigns to $f \oplus_{1/2} g$ cannot be because it is riskier in the sense of being a mean-preserving spread of f or g.

What other features could justify assigning a lower instrumental value to $f \oplus_{1/2} g$ than to f and to g? One possibly relevant difference is that, whereas f and g both have *two* potential outcomes with positive probability, $f \oplus_{1/2} g$ has *four*. But, clearly, having more potential outcomes doesn't *in itself* make a gamble riskier than another.[9] Another possibly relevant difference regards each gamble's *range*: the difference in value between its best and worst outcomes. Both f and g have a range of 2, while $f \oplus_{1/2} g$'s range is 3. But, again, having a wider range doesn't *in itself* make one gamble riskier than another.[10]

---

[8] Mean-preserving spreads were first explored in Rothschild and Stiglitz (1970), whose major contribution involved proving that the following three properties are equivalent: (1) $EU(X) \succeq EU(Y)$, for all concave utility-functions; (2) Y has more weight in its tails than X; (3) The outcomes of Y are distributed just like X's plus noise.

[9] Having more potential outcomes is consistent with being *less* risky. Consider, for example, the following gambles with the same mean. The first has three potential outcomes: a 99% chance of \$50, a 0.5% chance of \$49.99, and a 0.5% chance of \$50.01. The other has only two: a 50% chance of \$0 and a 50% chance of \$100. The former is clearly less risky than the latter despite having more potential outcomes.

[10] Given its insensitivity to the gambles' *probabilities*, range is at best a very limited indicator of riskiness. A gamble with a probability distribution concentrated around its mean might, in virtue of having some very low probability outliers, have a wider range—but be less risky—than one with a probability distribution concentrated in its tails.

4 Tie-breaking

Let's return to our story of Eunice and Reu. They are both indifferent between gamble *f* (betting *for* Easy Street) and gamble *g* (betting *'gainst* Easy Street), and must decide which to select. For Eunice, it's permissible to employ a tie-breaking procedure, like flipping a fair coin, in making her selection. For Reu, however, it is not.

Because Reu is indifferent between *f* and *g*, her reasons for taking the one are perfectly in balance with her reasons for taking the other. She has no rational basis for *choosing* the one over the other. This is—to borrow a distinction from Ullmann-Margalit and Morgenbesser (1977)—a situation that calls for *picking* rather than *choosing*. When we choose one option over another, we do so on the basis of the reasons we have that favor the one over the other. One way to think about picking, on the other hand, is that "when we are in a genuine picking situation we are in a sense transformed into a chance device that functions at random and effects arbitrary selections [...]" (Ullmann-Margalit and Morgenbesser, 1977, p. 773).

Suppose that Reu knows that, later, she will be in a picking situation: she'll have to select between options that she is (and knows she will continue to be) indifferent between. How should Reu value ending up in such a situation?

More concretely, let's suppose that Reu is in a hallway facing two doors. She knows that behind Door #1, there is a table on which sits three prizes: a ticket for gamble *f*, a ticket for gamble *g*, and $\$f$ $(= \$g)$, where $\$f$ $(\$g)$ is gamble *f*'s (*g*'s) sure-thing cash equivalent. If she opens Door #1, she gets to pick one of those three prizes. On the other hand, she knows that behind Door #2, there is a table with one prize on it: $\$2^1/_{16}$. If she opens Door #2, that's what she'll get.

It's (mostly) clear what Eunice would do in this situation: she would, first, open Door #1 and then select one of the three options (perhaps with the aid of some tie-breaking procedure). Even though Eunice might not know which prize she would select were she to open Door #1, doing so has greater value for her than opening Door #2. This is because, no matter how Eunice thinks about going on to pick between the three prizes if she opens Door #1, the value she assigns to opening Door #1—its expected value—is equal to the value she assigns to each of the three prizes behind it. And that value is greater than that of the $\$2^1/_{16}$ she is sure to get by opening Door #2 instead.

Opening Door #1 is a sensible thing for Eunice to do because—while opting for Door #1 doesn't *guarantee* a better outcome than opting for Door #2—she prefers each of the prizes behind Door #1 to the one behind Door #2. This is an example of Eunice obeying the following principle:

Menu Superiority

> If you know that φing will present you with a menu of options such that, for each of the items on that menu, you prefer it to all the items on the menu you'd be presented with if you didn't φ, then you rationally ought to prefer φing.

For example, when choosing between two restaurants, if you know that you prefer each of the dishes on offer at the first to all the dishes on offer at the second, it would be irrational to prefer going to the second. Eunice obeys this principle.

What about her sister, Reu? Does *she* obey the principle? Which does she prefer: Door #1 or Door #2? As I'll demonstrate in this section, it's not straightforward. Let's explore some of the possible, conflicting ways of approaching the question.

(1) Pick One of the Three Outright. Here is one way of thinking about Reu's situation. Because Reu is indifferent between the three options behind Door #1, for all she knows, she might pick any of the three. She has no reason to think she's more likely to pick any one of them than the others, so she should assign equal credence to each of the three possibilities.[11]

Gamble Between f and g and \$f

	Pick f		Pick g		Pick $\$f$	
	E	$\neg E$	E	$\neg E$	E	$\neg E$
$\oplus_{1/3}(f, g, \$f)$	4	2	1	3	$2^1/8$	$2^1/8$

[11] The reasoning here appears to appeal to something like the *Principle of Indifference*, which is controversial (see, for example, Hajek, 2003, p. 187-188). I don't think the argument turns on the truth of this principle in its full generality, however. For my purposes, it's enough that it be reasonable for someone like Reu to assign credences to her future actions in the way described above. It needn't be the case that she *must*—on the pains of irrationality—do so, only that this is an epistemically reasonable reaction to her situation. Given that Reu has no more reason to think she'll pick any of the options over any of the others, it isn't unreasonable for her to distribute her credence uniformly.

That said, one could argue that, in a case like this, (i) Reu is *radically uncertain* about what she might go on to pick and (ii) the uniquely epistemically rational response to radical uncertainty is to adopt *imprecise* probabilities over the various possibilities. (Seidenfeld (1988b, p. 310-311), discussing a closely related problem, advocates representing "this uncertainty with a (maximal) convex set of personal probabilities" over the admissible options. However, in Seidenfeld's example, the agent is unsure what she will go on to choose because she lacks a preference between her future options, not because she is indifferent between them. Steele (2010) discusses this approach as it applies to cases of indifference as well.) Generalizing a decision theory to handle imprecise probabilities is no easy task. Exploring whether this is a promising strategy for proponents of REUT to pursue is beyond the scope of this paper.

$$
\begin{aligned}
REU\left(\oplus_{1/3}(f, g, \$f)\right) =& 1 + \left(\frac{11}{12}\right)^2 \cdot (2 - 1) + \left(\frac{8}{12}\right)^2 \cdot (2^1/_8 - 2) \\
&+ \left(\frac{4}{12}\right)^2 \cdot (3 - 2^1/_8) + \left(\frac{1}{12}\right)^2 \cdot (4 - 3) \\
=& 1 + \frac{121}{144} + \frac{8}{144} + \frac{14}{144} + \frac{1}{144} = 2
\end{aligned}
$$

If Reu thinks that, were she to open Door #1, it's equally likely that she'd select any one of the three prizes over the others, then opening Door #1 is like a lottery with a $1/3$ chance of paying out gamble *f*, a $1/3$ chance of paying out gamble *g*, and a $1/3$ chance of paying out \$*f* (which, for Reu, is \$2¹/8). Given Reu's attitude toward risk, such a lottery is valued at \$2. Because opening Door #2 will result in \$2¹/16 for sure, doing so has higher value for her than opening Door #1. So, if this is the procedure that Reu knows she'll use when picking between the three prizes behind Door #1, she should open Door #2 instead.

But, by preferring Door #2 to Door #1, Reu violates *Menu Superiority*: she knows, for each of the prizes behind Door #1, that she prefers it to the prize behind Door #2 and yet prefers opening Door #2. In fact, given that one of the things she could do is to open Door #1 and then select the sure-thing \$2¹/8, there is something Reu could do that would *guarantee* a better outcome than the one she knows she'll bring about by opening Door #2. And yet—if it's reasonable for Reu to think that she will employ such a procedure when selecting between the prizes behind Door #1—she should nevertheless prefer to open Door #2.

That might seem absurd. But there are other ways we might approach Reu's situation—ways which might prove to be less absurd. Let's take a look.

(2) Use a Sequential Pairwise Procedure. Alternatively, when picking between the three prizes behind Door #1, Reu might employ a procedure that involves evaluating the options pair-by-pair. There are two such procedures: the first, let's call *Process of Elimination*; and the second, let's call the *Tournament Method*. Let's look at each in turn.

- *Process of Elimination.* When deciding between three options, $\{X, Y, Z\}$, you first decide whether to select one of the options outright or to eliminate it from the running; if you opt to eliminate it, then you go on to decide between the remaining two.

Process of Elimination (Z)

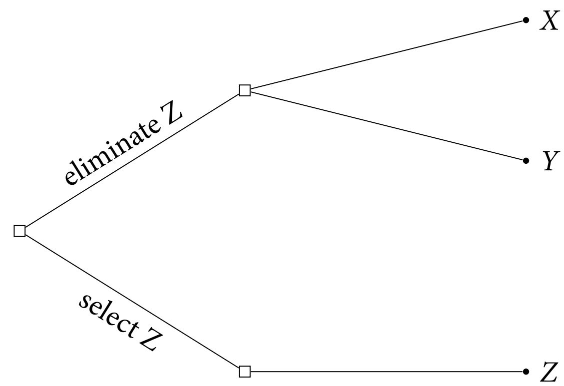

Because there are three ways to choose one from a set of three, there are three different ways of employing this procedure: you can first decide whether to select or eliminate X, choosing between Y and Z if you opt for elimination (call this *Eliminate X*); or you can first decide whether to select or eliminate Y, choosing between Z and X if you opt for elimination (*Eliminate Y*); or, lastly, you can decide whether to select or eliminate Z, choosing between X or Y if you opt for elimination (*Eliminate Z*). Let's call these *ways of setting the agenda*. In this case, there are three ways of setting the agenda.

- *Tournament Method.* When deciding between three options, $\{X, Y, Z\}$, you first decide which out of two of them will proceed to the next round; then, at the next round, the "winner" of the first round faces-off against the remaining option.

Tournament Method (Z Gets a "Bye")

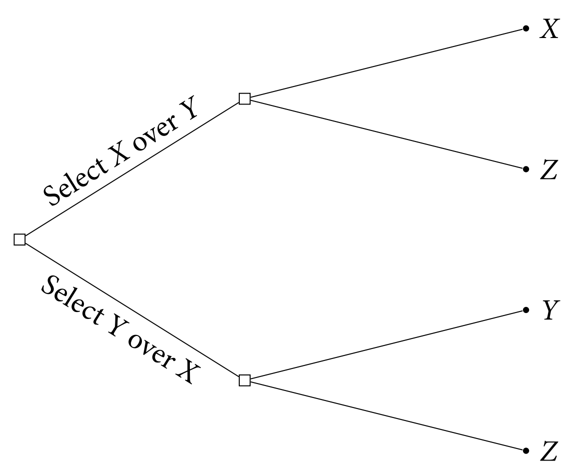

Again, because there are three options, there are three different ways of employing this procedure: one where X gets a bye, one where Y gets a bye, and one where Z gets a bye. Here, too, there are three ways of setting the agenda.

We have two different sequential pairwise procedures and, for each, three ways of setting the agenda. When picking between the three prizes between Door #1, does it matter which procedure is used and how the agenda is set? Should it? If you're Eunice, it doesn't matter. No matter which procedure she thinks she might implement and however she might set the agenda were she to open Door #1, opening Door #1 will have greater expected value than opening Door #2. Furthermore, because Eunice is indifferent between the three prizes behind Door #1, none of the procedures—and none of the ways of setting the agenda—make it any more or less likely that she'll end up with one of the prizes rather than another. Consequently, she will have no reason to favor adopting any one of these procedures over any other. Just as she is indifferent between the results of the procedure—that is, the three prizes on the table—she is likewise indifferent between the procedures themselves. Eunice obeys the following principle:

> Agenda Invariance
>
> When deciding from a menu of options, how you set the agenda should have no effect on what it's rational to do.[12]

When choosing an entree from a menu, it shouldn't matter whether you start at the top or the bottom of the list. If it would be irrational for you to choose the Jellied Eel were you to start at the top, it should be irrational for you to choose the Jellied Eel were you to start from the bottom instead. How the decision is "framed" shouldn't matter.[13]

But as I'll demonstrate, if you're Reu, it *does* matter. Reu, unlike her sister, violates *Agenda Invariance*: for her, how the agenda is set does effect what it's rational for her to do. Let's consider each of the two sequential pairwise procedures outlined above, in turn.

- **Process of Elimination.** As mentioned before, when there are three options on the menu, there are three ways of setting the agenda. In Reu's case, the menu behind Door #1 includes the ticket for gamble *f*, the ticket for

[12] More precisely: When deciding from a menu M of options, if there is a way of setting the agenda $A(M)$ such that, were you to set the agenda that way, rational choice would result in the selection of $X \in M$, then *every* way of setting the agenda is likewise such that, were you to use *it*, rational choice would result in the selection of X. What it's rational for you to do is *invariant* across agendas.

[13] In a series of influential papers (Tversky and Kahneman, 1979, 1984, 1986), Daniel Kahneman and Amos Tversky argue that, for many of us, how a decision-problem is framed *does* matter. However, their examples pertain to how particular *options* are framed—that is, how various features of an option are described—and not to the way in which those options are considered. Holding fixed how the options themselves are described, should it matter e.g. the order in which they are considered? Furthermore, Kahneman and Tversky are engaged in a descriptive project, not a normative one. They agree that principles, like the one above, are "normatively essential" (Tversky and Kahneman, 1986).

gamble *g*, and \$*f* (= \$2 1/8)—the sure-thing cash equivalent. Thus, there are a three corresponding ways of setting the agenda when using *Process of Elimination*: she can decide whether or not to eliminate *f*, or she can decide whether or not to eliminate *g*, or she can decide whether or not to eliminate \$*f*. Let's explore each of these.

Eliminate f. On this way of setting the agenda, Reu first decides whether to select or to eliminate gamble *f*. If she opts to eliminate *f*, she's then faced with a choice between gamble *g* and \$2 1/8—its sure-thing cash equivalent.

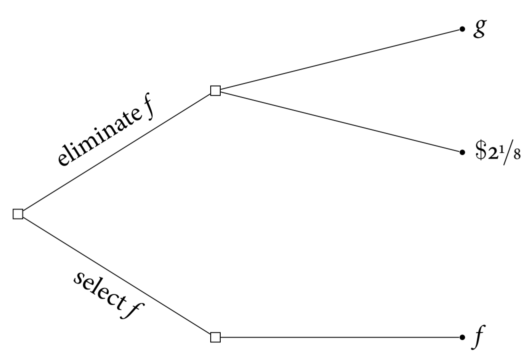

What should Reu do? She knows that if she elects to *eliminate f*, she will face a choice between two prizes—gamble *g* and \$2 1/8—between which she is indifferent. Because she is indifferent between the two, for all she knows, she might pick either when facing such a choice. So, she assigns equal credence to ending up with either of the prizes in the event that she eliminates *f*. Thus, *eliminate f* is like a lottery with a 1/2 chance of paying out gamble *g* and a 1/2 chance of paying out \$2 1/8. Given Reu's attitude toward risk, such a lottery is valued at \$1 63/64.[14] Because gamble *f*—which is obviously what will result were Reu to elect to *select f* rather than to eliminate it—is valued at \$2 1/8, Reu should prefer selecting gamble *f* over eliminating it.

Result: If this is the way the agenda is set, Reu will select gamble *f*.

Eliminate g. On this way of setting the agenda, Reu first decides whether to select or to eliminate gamble *g*. If she opts to eliminate *g*, she's then faced with a choice between gamble *f* and \$2 1/8—its sure-thing cash equivalent.

[14] This is not obvious. See appendix B for details of the calculation.

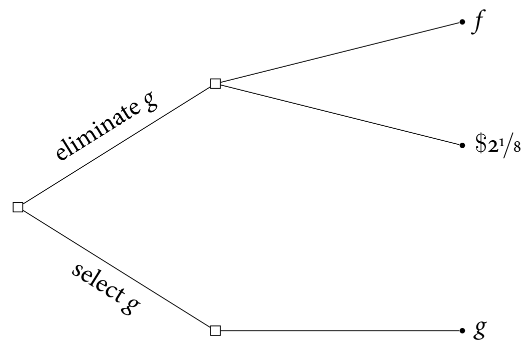

Because Reu is indifferent between gamble *f* and $2 1/8, the option to *eliminate g* is like a 50/50 lottery between the two, which is valued at $2 5/64 given her attitude toward risk.[15] Because gamble *g*—which is what will result were she to *select g* rather than eliminate it—is valued at $2 1/8, Reu should prefer selecting gamble *g* over eliminating it.

Result: If this is the way the agenda is set, Reu will select gamble *g*.

Eliminate $2 1/8. On this way of setting the agenda, Reu decides whether to select or to eliminate $2 1/8. If she opts to eliminate it, she goes on to face a choice between the two gambles.

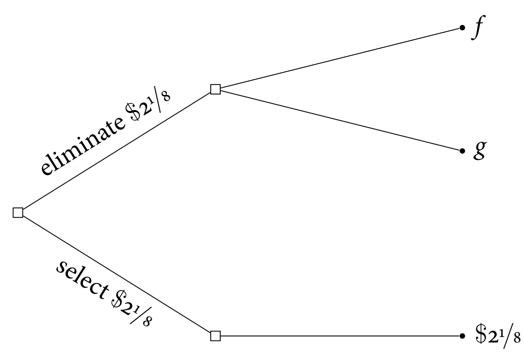

Because Reu is indifferent between the two gambles, choosing between them is like a 50/50 lottery. As we saw in Section 3, a 50/50 lottery between gamble *f* and gamble *g* is valued at $2 1/32. This is worth less than the $2 1/8 Reu will get by selecting it outright. So, that's what she should do.

Result: If this is the way the agenda is set, Reu will select the $2 1/8.

There are three things to note. First, if Reu uses *Process of Elimination* to select between the three prizes behind Door #1, her indifference among the three is recapitulated in her decision about how to set the agenda. Each way of setting the agenda results in each one of the prizes. Second, because each

[15] Again, see appendix B for the calculations.

of the three prizes is the result of one of the ways of setting the agenda, Reu violates *Agenda Invariance*. Different ways of setting the agenda result in different prizes, so the agenda *does* have an effect on what it's (seemingly) rational for Reu to do.

Finally, what should Reu do—open Door #1 or Door #2—if she knows that, were she to open Door #1, she would use *Process of Elimination* to select between the three prizes? The answer depends on how Reu thinks she might set the agenda were she to open Door #1. Because each way of setting the agenda will result in a prize that is valued at $\$2^1/_8$, whereas the prize behind Door #2 is only $\$2^1/_{16}$, if Reu is certain of how she would set the agenda, she should prefer Door #1 to Door #2. However, Reu knows that, just as she is indifferent between the three prizes themselves, she is indifferent between the three ways of setting the agenda—and, so, just as it's rational for her to be unsure of which of the three prizes she'd select, she shouldn't be sure of how she'd set the agenda. If, for instance, she thinks she is equally likely to set the agenda in any of the three ways, then (echoing the argument from above regarding the prizes themselves) opening Door #1 is like a lottery with an equal chance of paying out any of the three prizes, which (as we saw above) she values at \$2—and, thus, she should prefer Door #2. But it isn't obvious that she should think she is equally likely to set the agenda in any of the three ways—for, just as there are various ways she might decide between the three prizes themselves, there are various ways she might decide between the three ways of setting the agenda. And, as we shall see, there are ways of making this meta-decision (i.e., the decision about how to set the agenda for deciding which of the prizes to select) that favor some ways of setting the agenda over others.

- **Tournament Method.** With this method, as well, there are three different ways of setting the agenda. Reu can decide to allow gamble *f* to get a bye into the next round, or to allow gamble *g* to get a bye into the next round, or to allow the $\$2^1/_8$ sure-thing cash equivalent to get a bye into the next round. Let's look at each.

    f gets a bye. On this way of setting the agenda, Reu first decides whether or not to select gamble *g* over $\$2^1/_8$. Then, whichever is selected is put in contention with gamble *f*. And Reu decides which of those two to take home.

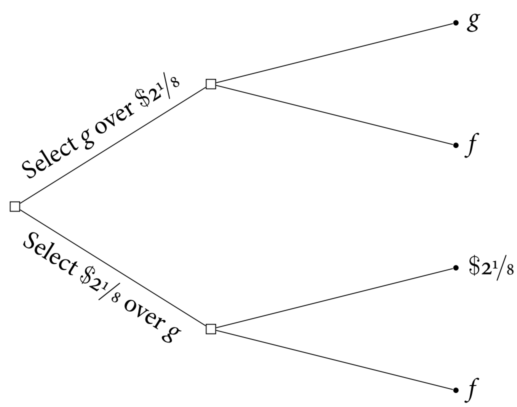

Reu knows that if she selects gamble *g* over $\$2\frac{1}{8}$, she'll then face a choice between gamble *g* and gamble *f*. Because she is indifferent between the two, for all she knows, she might pick either one. So, selecting gamble *g* over $\$2\frac{1}{8}$ is like a 50/50 lottery between gamble *g* and gamble *f*. Given her attitude toward risk, Reu values such a lottery at $\$2\frac{1}{32}$. On the other hand, if she selects $\$2\frac{1}{8}$ over gamble *g*, she'll then face a choice between $\$2\frac{1}{8}$ and gamble *f*. For similar reasons, she can think of this choice as akin to a 50/50 lottery between the two. Given her attitude toward risk, she values it at $\$2\frac{5}{64}$.[16] The latter is better than the former, so Reu should opt to select $\$2\frac{1}{8}$ over gamble *g*.

Result: If this is the way the agenda is set, Reu will take home either $\$2\frac{1}{8}$ or gamble *f*. And this way of setting the agenda is worth $\$2\frac{5}{64}$ to her.

g gets a bye. On this way of setting the agenda, Reu decides whether or not to select gamble *f* over $\$2\frac{1}{8}$. Whichever wins goes on to face gamble *g* in the final round.

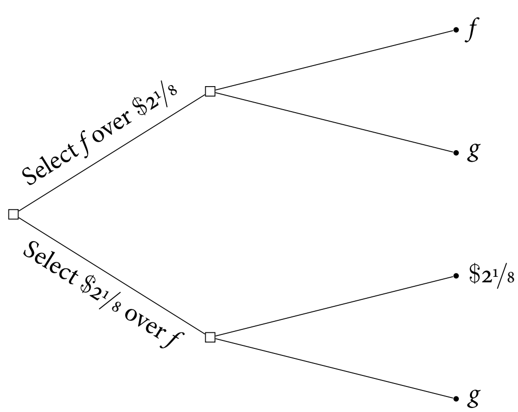

For reasons analogous to those in the previous example, we can represent Reu's choice as between two 50/50 lotteries—one between gam-

[16] See appendix B for details.

ble *f* and gamble *g*, and the other between \$2¹⁄8 and gamble *g*. The former, as we saw, is valued at \$2¹⁄32, while the latter is valued at \$1⁶³⁄64.[17] The former is better than the latter, so Reu should opt to select gamble *f* over \$2¹⁄8.

Result: If this is the way the agenda is set, Reu will take home either gamble *f* or gamble *g*. And this way of setting the agenda is worth \$2¹⁄32 to her.

\$2¹⁄8 gets a bye. Lastly, on this way of setting the agenda, Reu first decides whether or not to select gamble *f* over *g*. The winner faces off against \$2¹⁄8—their sure-thing cash equivalent.

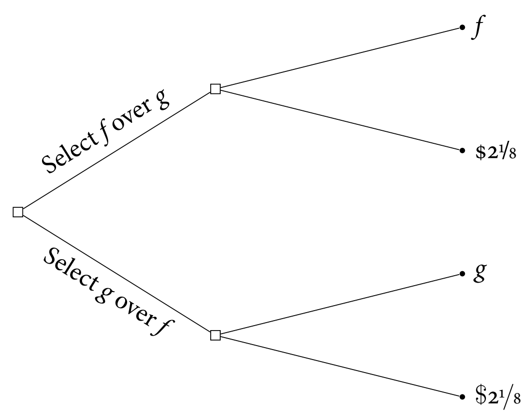

Like above, we can compare the 50/50 lottery between gamble *f* and \$2¹⁄8 with the 50/50 lottery between gamble *g* and \$2¹⁄8. The former is valued at \$2⁵⁄64, while the latter is valued at \$1⁶³⁄64. The former is better than the latter, so Reu should opt to select gamble *f* over *g*.

Result: If this is the way the agenda is set, Reu will take home either gamble *f* or \$2¹⁄8. And this way of setting the agenda is worth \$2⁵⁄64 to her.

Again, there are three things to note. First, unlike with *Process of Elimination*, Reu values these different ways of setting the agenda differently—she disprefers the agenda in which gamble *g* gets a bye to the other two. She values picking from the menu $\{f, \$2^1/_8\}$ over picking from the menu $\{f, g\}$. Setting the agenda in a way that she prefers, though, effectively removes gamble *g* from the running. (This is odd because, given Reu's indifference between gamble *g* and the other prizes, it seems like it very much *is* still in the running.) Second, we have yet another example of Reu violating *Agenda Invariance*. Not only might different agendas result in different prizes, this is a case in which Reu has a clear preference for some of

[17] See appendix B for details.

> the ways of setting the agenda over others. And, lastly, if Reu knew that, were she to open Door #1, she'd select between the prizes using the *Tournament Method*, then, because she prospectively values deciding in this way at $\$2^5/_{64}$, she'd prefer to open Door #1 over Door #2.

Should Reu open Door #1 or Door #2? It seems like it depends on how she believes she might go about deciding between the prizes between Door #1. If she thinks she'll implement the *Tournament Method*, then she should open Door #1. (Which prize she'll end up with will depend on how she sets the agenda.) If, on the other hand, she thinks she'll just pick one of three outright, then she should open Door #2 instead. It's less clear what she should do if she thinks she'll implement *Process of Elimination* because that procedure just recapitulates the first-order decision at the level of agendas—and so which door she should open depends on which procedure she thinks she might use to select which of the ways to set the agenda. If she thinks she'll just pick one of three agendas at random, then she should open Door #2; however, if she thinks she'll use the *Tournament Method* to select between the agendas, she should ultimately open Door #1.[18] How she thinks her future decision will be structured—a decision only concerning prizes that she is indifferent between—makes a (some might say, "implausible") difference to how she ought to act.

5 Lessons, Objections, and Replies

Because the value of picking from the menu $\{f, g, \$2^1/_8\}$ depends on the procedure that's used, it lacks a stable value. But there should be a univocal answer about whether rationality requires Reu to pick from that menu (which is what's behind Door #1) or accept the $\$2^1/_{16}$ (which is what's behind Door #2). And if no stable value can be assigned to picking from the menu, there will not be a univocal answer. Therefore, we should reject REUT.[19]

That's, very roughly, the argument I've raised against REUT.[20] Is it right? Let's look at some objections and replies.

[18] What if she thinks she'll use *Process of Elimination* to determine which of the ways to set the agenda? This, like before, recapitulates the previous decision: there are three different ways to set the agenda for making *that* decision, each one resulting in a different one of the ways of setting the agenda for the first-order decision, and each of which is equally valuable to her. A regress looms.

[19] Note, however, that there is no analogous problem for EUT. Although Eunice might consider using the same procedures as her sister (e.g., picking one of the three prizes at random, using *Process of Elimination*, using the *Tournament Method*), each will have the same value for her as each of the prizes themselves. So, opening Door #1, for Eunice, has a stable value: it's the same value she assigns to each of the pries she knows awaits if she does.

[20] The argument presented above shares certain similarities with the one Seidenfeld (1988a) develops against decision theories that jettison the Independence axiom. REUT, in virtue of violating Independence, is subject to Seidenfeld's objection. The argument is developed and further refined in

Response 1: It's not true that picking from the menu $\{f, g, \$2^1/_8\}$ lacks a stable value. That doesn't follow from the fact that its value depends on the procedure that will be (or that Reu *believes* will be) used.

1.1 The value of picking from the menu is given, externally, by how the decision-tree looks. Different decision-trees correspond to different problems. So there's nothing objectionable about violating *Agenda Invariance* because there's nothing objectionable about treating *different* decision problems differently.

Response to 1.1: The decision-tree is merely a representation—one that, while subject to external constraints, isn't fully determined by your external circumstances. We aren't imagining that Reu is literally facing external choice-points—like actual forks in an actual road. Rather, the different decision-trees are meant to represent different ways she might structure her internal *deliberation* among the options.

1.2 Just as Reu might be uncertain about which of the items will be picked from the menu were she faced with having to select an item from it (call this a *first-order decision*), Reu should also be uncertain about which of the procedures she will employ (call this a *second-order decision*). Likewise, if one's uncertain about which of the procedures one will employ—if, for example, the value of using a procedure depends on how the agenda is set—one should also be uncertain about *that*. And so on and so forth. The value of the picking situation can be found by iterating this process until it terminates in some stable value.

Response to 1.2: That's true only if a stable evaluation emerges. But, in this case, it's not clear that one does. As we've seen, when selecting between three items, there are seven possibilities: Reu can select one of the three at random, use one of the three versions of *Process of Elimination*, or one of the three versions of the *Tournament Method*. Out of these possibilities, the three versions of *Process of Elimination* have the highest risk-weighted expected utility.[21] If Reu expects herself to employ a selection procedure

(Seidenfeld, 1988b, 2000a,b) in response to criticism by McClennen (1988) and Rabinowicz (1995, 1997, 2000). (See Steele, 2010, for further discussion of the dialectic.) My argument differs from Seidenfeld's in some key respects. His argument relies on a dynamic coherence constraint that requires one's evaluations of plans to be unchanged by substitutions, at choice nodes, of indifferent options. He shows that decision theories that jettison Independence violate this constraint—they prescribe evaluating plans in a way that is, according to Seidenfeld, dynamically incoherent.

[21] Recall, Reu values each version of *Process of Elimination* at $\$2^1/_8$; she values picking one of the three randomly at \$2; she values two of the versions of the *Tournament Method* at $\$2^5/_{64}$, and the other at $\$2^1/_{32}$.

only if there is no other she prefers to it, then she should expect herself to use some version of *Process of Elimination* in selecting the prize behind Door #1. But because the higher-order decision about which *version* of *Process of Elimination* to employ—the question of how to set the agenda—recapitulates the first-order decision about which prize to select, it's not clear that iterating *Process of Elimination* will result in a stable value assignment. At each level of application, the higher-order decision about how to set the agenda recapitulates the decision below it. No stable value emerges, only a regress.

Response 2: It's not a mark against REUT that there's no univocal answer about what rationality requires of you in these situations. Sometimes, there's just no fact of the matter about what you are rationally required to do.

Response to 2: This might not be so bad a response if the cases in which there is no fact of the matter about what you are rationally required to do are relatively rare or outlandish. But that's not the case here. Whenever an agent is indifferent between risky gambles, it will be possible to construct—or just wander one's way into—a situation like this one. This is, then, at best, a position of last resort.

Response 3: The value of picking from the menu $\{f, g, \$2\frac{1}{8}\}$ does *not* depend on the procedure that's used. You shouldn't treat your future decisions as something to be predicted, but rather as something to be *decided:* pick the *plan* that you like best (where a 'plan' is a complete path through a decision-tree). In this case, you're indifferent between all of them.

Response to 3: This is to endorse, what is sometimes called, Resolute Choice.[22] One should settle on the overall plan that one most desires, and then stick to it—even if that involves acting counter to your preferences in the future. However, if REUT avails itself of this response, then, as Thoma (2019) convincingly argues, its recommendations will be approximately the same as expected utility theory's. This move, then, is at best a Pyrrhic victory. It rescues REUT from objection, but in so doing bleeds it of its distinctive content.

6 Conclusion

I've argued that REUT has some counterintuitive consequences regarding cases of tie-breaking. It doesn't offer clear recommendations about what is rationally required of you when deciding between options you're indifferent between. REUT

[22] See Machina (1989) and McClennen (1990, 1997) for classic presentations of the view.

is meant to be an all-encompassing theory of rational action. This should include decisions in which a source of uncertainty is your own future indecisiveness. REUT, however, cannot handle all such cases in plausible way—it appears to violate either *Menu Superiority* or *Agenda Invariance* or both.

Furthermore, I've raised some worries that REUT doesn't offer an appropriate measure of riskiness to begin with. So, while it might be true that EUT fails to adequately account for differing attitudes toward risk, REUT isn't the appropriate remedy. I have no suggestion for what would be an appropriate remedy. But let's not rue a cure that's worse than the disease.

A Variance and Other Measures of Risk

In §3, I claimed that the gambles f, g, and $f \oplus_{1/2} g$ all have the same *variance*: 0.75. I'll now show why that is the case.

Recall that, for any gamble $h = \{x_1, E_1; x_2, E_2; \ldots x_n, E_n\}$ (where $EU(h)$ is its expected value), the variance of h is:

$$\textsc{Var}(h) = \sum_{i=1}^{n} Cr(E_i) \cdot (u(x_i) - EU(h))^2$$

Therefore,

$$\begin{aligned}\textsc{Var}(f) &= 1/4 \cdot (4 - 2.5)^2 + 3/4 \cdot (2 - 2.5)^2 \\ &= 1/4 \cdot (1.5)^2 + 3/4 \cdot (-.5)^2 \\ &= 1/4 \cdot (2.25) + 3/4 \cdot (.25) \\ &= .5625 + .1875 = .75\end{aligned}$$

$$\begin{aligned}\textsc{Var}(g) &= 1/4 \cdot (1 - 2.5)^2 + 3/4 \cdot (3 - 2.5)^2 \\ &= 1/4 \cdot (-1.5)^2 + 3/4 \cdot (.5)^2 \\ &= 1/4 \cdot (2.25) + 3/4 \cdot (.25) \\ &= .5625 + .1875 = .75\end{aligned}$$

$$\begin{aligned}
\textsc{Var}(f \oplus_{1/2} g) &= {}^1\!/_8 \cdot (4 - 2.5)^2 + {}^3\!/_8 \cdot (3 - 2.5)^2 + {}^3\!/_8 \cdot (2 - 2.5)^2 + {}^1\!/_8 \cdot (1 - 2.5)^2 \\
&= {}^1\!/_8 \cdot (1.5)^2 + {}^3\!/_8 \cdot (.5)^2 + {}^3\!/_8 \cdot (-.5)^2 + {}^1\!/_8 \cdot (-1.5)^2 \\
&= {}^1\!/_8 \cdot (2.25) + {}^3\!/_8 \cdot (.25) + {}^3\!/_8 \cdot (.25) + {}^1\!/_8 \cdot (2.25) \\
&= {}^1\!/_4 \cdot (2.25) + {}^3\!/_4 \cdot (.25) \\
&= .5625 + .1875 = .75
\end{aligned}$$

The gambles f and g have the same variance, as does their 50/50 mixture. This is no coincidence. In general, if two gambles have the same mean and variance, then any probabilistic mixture of the two will also have the same mean and variance.

A related measure of the "dispersion" of a gamble's outcomes is its *mean absolute deviation*:

$$\text{MAD}(h) = \sum_{i=1}^{n} Cr\,(E_i) \cdot |u(x_i) - EU(h)|$$

But, as inspecting the previous calculations might make clear, this change makes no difference to the three gambles' relative standing: $\text{MAD}(f) = \text{MAD}(g) = \text{MAD}(f \oplus_{1/2} g)$.[23]

One criticism of both of these measures is that they treat correspondingly large deviations from the mean the same whether those deviations are positive or negative. However, variation *below* the mean might be thought to contribute more to a gamble's riskiness than correspondingly large variation above the mean; downside risk, it might be thought, weighs more heavily than upside risk. If that's right, we might might want to exclusively focus on a gamble's *lower semi-variance*: the probabilistically-weighted squared deviation from the mean of those outcomes whose values are below the mean. Or, at least, we might want to weight a gamble's lower semi-variance ($\textsc{SemiVar}^-$) more heavily than its upper semi-variance ($\textsc{SemiVar}^+$) when assessing its riskiness.

But, as another glance at the previous calculations can reveal, there are no weights α_1 and α_2 such that

$$\begin{aligned}
\alpha_1 \cdot \textsc{SemiVar}^-(f) + \alpha_2 \cdot \textsc{SemiVar}^+(f) &= \alpha_1 \cdot \textsc{SemiVar}^-(g) + \alpha_2 \cdot \textsc{SemiVar}^+(g) \\
&< \alpha_1 \cdot \textsc{SemiVar}^-(f \oplus_{1/2} g) + \alpha_2 \cdot \textsc{SemiVar}^+(f \oplus_{1/2} g)
\end{aligned}$$

This is because

[23] Somewhat surprisingly, their mean absolute deviations also equal 0.75. That *is* a coincidence. Typically, a gamble's mean absolute deviation and its variance will *not* be the same.

$$\begin{aligned}\textsc{SemiVar}^-(f) &= {}^3\!/_4 \cdot (-.5)^2 \\ &= .1875 \\ \textsc{SemiVar}^+(f) &= {}^1\!/_4 \cdot (1.5)^2 \\ &= .5625\end{aligned}$$

$$\begin{aligned}\textsc{SemiVar}^-(g) &= {}^1\!/_4 \cdot (-1.5)^2 \\ &= .5625 \\ \textsc{SemiVar}^+(g) &= {}^3\!/_4 \cdot (.5)^2 \\ &= .1875\end{aligned}$$

$$\begin{aligned}\textsc{SemiVar}^-(f \oplus_{1/2} g) &= {}^3\!/_8 \cdot (-.5)^2 + {}^1\!/_8 \cdot (-1.5)^2 \\ &= .375 \\ \textsc{SemiVar}^+(f \oplus_{1/2} g) &= {}^1\!/_8 \cdot (1.5)^2 + {}^3\!/_8 \cdot (.5)^2 \\ &= .375\end{aligned}$$

And so the only way to weight *f*'s and *g*'s lower and upper semi-variances so that their sums are equal is if they are weighed equally: $\alpha_1 = \alpha_2$. But in that case $f \oplus_{1/2} g$ will be ranked right alongside *f* and *g* again. Even if downside risk should loom larger than upside risk, $f \oplus_{1/2} g$ isn't any riskier (in that sense) than *both f* and *g*.

B REU Calculations

Let's look at the risk-weighted expected value of deciding between each of the pairs. First, consider deciding between gamble *f* and its sure-thing cash equivalent \$*f*, which—because Reu has no reason to think she is more likely to select the one rather than the other—corresponds to a 50/50 lottery between the two.

Gamble Between f and \$f

	Pick *f*		Pick \$*f*	
	E	$\neg E$	E	$\neg E$
$f \oplus_{1/2} \$f$	4	2	2¹/8	2¹/8

$$\begin{aligned}REU\left(f \oplus_{\frac{1}{2}} \$f\right) &= 2 + \left(\frac{5}{8}\right)^2 \cdot (2\tfrac{1}{8} - 2) + \left(\frac{1}{8}\right)^2 \cdot (4 - 2{}^1\!/_8) \\ &= 2 + \frac{25}{512} + \frac{15}{512} = 2{}^5\!/_{64}\end{aligned}$$

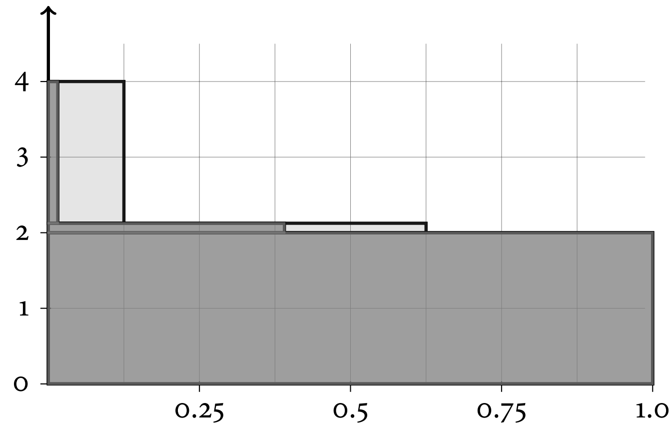

Figure 4: The risk-weighted expected utility of a fifty-fifty lottery between gamble f and its sure-thing utility equivalent $\$f$ (2.078125).

If Reu thinks it's equally likely, when picking between f and $\$f$, that she'll end up with either, then picking between the two is valued at $\$2\frac{5}{64}$. (See Figure 4.)

Next, consider the decision between gamble g and its sure-thing cash equivalent $\$g$ (which, again, because Reu is indifferent between f and g, is identical to $\$f$).

Gamble Between g and $g

	Pick g		Pick $\$g$	
	E	$\neg E$	E	$\neg E$
$g \oplus_{1/2} \$g$	1	3	$2\frac{1}{8}$	$2\frac{1}{8}$

$$\begin{aligned} REU\left(g \oplus_{1/2} \$g\right) &= 1 + \left(\frac{7}{8}\right)^2 \cdot \left(2\tfrac{1}{8} - 2\right) + \left(\frac{3}{8}\right)^2 \cdot \left(3 - 2\tfrac{1}{8}\right) \\ &= 1 + \frac{441}{512} + \frac{63}{512} = 1\tfrac{63}{64} \end{aligned}$$

If Reu thinks it's equally likely, when picking between g and $\$g$, that she'll end up with either, then picking between the two is valued at $\$1\frac{63}{64}$. (See Figure 5.)

Acknowledgements

For feedback and discussion, I would like to thank Luc Bovens, Richard Bradley, Geoff Brennan, Jennifer Carr, Tom Dougherty, Caspar Hare, Ze'ev Goldschmidt, Simone Gubler, Patricia Rich, Geoff Sayre-McCord, H. Orri Stefánsson, Johanna Thoma, Ittay Nissan-Rozen, Fabian Wendt, the audiences at the Mini-Workshop

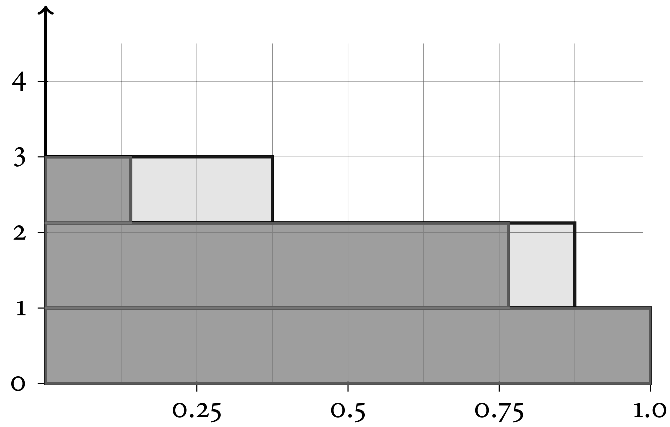

Figure 5: The risk-weighted expected utility of a fifty-fifty lottery between gamble *g* and its sure-thing utility equivalent $*g* (1.984375).

on Attitudes to Risk at the Hebrew University of Jerusalem, the UNC PPE retreat, and the University of San Diego. I'd also like to thank two anonymous reviewers whose generous comments helped improve the paper.

References

Karl Borch. A note on uncertainty and indifference curves. *The Review of Economic Studies*, 36:1–4, 1969. URL `https://doi.org/10.2307/2296336`.

Lara Buchak. *Risk and Rationality*. Oxford University Press, 2013.

James Dreier. Rational preference: Decision theory as a theory of practical rationality. *Theory and Decision*, 40:249–76, 1996.

Alan Hajek. Conditional probability is the very guide of life. In Jr. H. E. Kyburg and M. Thalos, editors, *Probability is the very guide of life: The philosophical uses of chance*, pages 187–88. Open Court, Chicago, 2003.

David Johnstone and Dennis Lindley. Mean-variance and expected utility: The borch paradox. *Statistical Science*, 28(2):223–237, 2013. URL `10.1214/12-STS408`.

Veronika Köbberling and Peter P. Wakker. Preference foundations for nonexpected utility: A generalized and simplified technique. *Mathematics of Operations Research*, 28(3):395–423, 2003.

Mark J. Machina and David Schmeidler. A more robust definition of subjective probability. *Econometrica*, 60(4):745–780, 1992.

M.J. Machina. Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, 27:1622–1668, 1989.

Harry Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952.

Harry Markowitz. *Portfolio Selection: Efficient Diversification of Investments*. Monograph 16. Cowels Foundation for Research in Economics at Yale University, 1959.

Edward F. McClennen. *Rationality and Dynamic Choice: Foundational Explorations*. Cambridge University Press, 1990.

Edward F. McClennen. Pragmatic rationality and rules. *Philosophy and Public Affairs*, 26(3):210–258, 1997.

E.F. McClennen. Ordering and independence: a comment on professor seidenfeld. *Economics and Philosophy*, 4:298–308, 1988.

Graham Oddie and Peter Milne. Act and value: Expectation and the representability of moral theories. *Theoria*, 57(1-2):42–76, 1991.

Wlodek Rabinowicz. To have one's cake and eat it, too: Sequential choice and expected-utility violations. *The Journal of Philosophy*, 92(11):586–620, November 1995.

Wlodek Rabinowicz. On seidenfeld's criticism of sophisticated violations of the independence axiom. *Theory and Decision*, 43:279–292, 1997.

Wlodek Rabinowicz. Preference stability and substitution of indifferents: A rejoinder to seidenfeld. *Theory and Decision*, 48:311–318, 2000.

Michael Rothschild and Joseph Stiglitz. Increasing risk: I. a definition. *Journal of Economic Theory*, 2(3):225–243, 1970.

Paul A. Samuelson. General proof that diversification pays. *The Journal of Financial and Quantitative Analysis*, 2(1):1–13, 1967.

Leonard J Savage. *The Foundations of Statistics*. Wiley, 1954.

Teddy Seidenfeld. Decision theory without 'independence' or without 'ordering'. *Economics and Philosophy*, 4(2):267–90, 1988a.

Teddy Seidenfeld. Rejoinder [to Hammond and McClennen]. *Economics and Philosophy*, 4(2):309–315, 1988b.

Teddy Seidenfeld. Substitution of indifferent options at choice nodes and admissibility: A reply to Rabinowicz. *Theory and Decision*, 48:319–322, 2000a.

Teddy Seidenfeld. The independence postulate, hypothetical and called-off acts: A further reply to Rabinowicz. *Theory and Decision*, 48(2):319–322, 2000b.

Katie Siobhan Steele. What are the minimal requirements of rational choice? arguments from the sequential-decision setting. *Theory*, 68:463–487, December 2010.

Johanna Thoma. Risk aversion and the long run. *Ethics*, 129(2):230–253, 2019.

Amos Tversky and Daniel Kahneman. Prospect theory: An analysis of decision under risk. *Econometrica*, 47:263–291, 1979.

Amos Tversky and Daniel Kahneman. Choice, values, and frames. *American Psychologist*, 39:341–350, 1984.

Amos Tversky and Daniel Kahneman. Rational choice and the framing of decisions. *The Journal of Business*, 59(4 Part 2: The Behavioral Foundations of Economic Theory):S251–S278, 1986.

Edna Ullmann-Margalit and Sidney Morgenbesser. Picking and choosing. *Social Research*, 44(4):757–785, 1977.