Making Ado Without Expectations

Mark Colyvan
Department of Philosophy, University of Sydney
mark.colyvan@sydney.edu.au

Alan Hájek
School of Philosophy, Australian National University
alan.hajek@anu.edu.au

This paper is a response to Paul Bartha's 'Making Do Without Expectations'. We provide an assessment of the strengths and limitations of two notable extensions of standard decision theory: relative expectation theory (RET) and Paul Bartha's relative utility theory (RUT). These extensions are designed to provide intuitive answers to some well-known problems in decision theory involving gaps in expectations. We argue that both RET and RUT go some way towards providing solutions to the problems in question but neither extension solves all the relevant problems.

1. Expectation gaps

Once upon a time, all was well with expected utility theory — or so it seemed. The 17th century authors of the *Port Royale Logic* (Arnauld and Nicole 1964) already regarded rational decision as being a matter of maximizing expected utility: a probability-weighted average of the utilities associated with each combination of action and state of the world. Ramsey (1931), von Neumann and Morgenstern (1944), Savage (1954), and Jeffrey (1983) gave foundational support to this idea. Then along came Newcomb's problem (Nozick 1969), which prompted various formulations of causal decision theory (e.g. Gibbard and Harper 1981; Lewis 1981; Joyce 1999). Yet causal decision theorists still advocate maximizing expected utility; they just offer different proposals for what the probability weights should be. In this sense, disputes among these various protagonists were and are really *intramural*, taking place against a backdrop of broad agreement on the general framework of expected utility theory.

However, that theory's health has faced more serious threats. An early bug was the St. Petersburg game, which has infinite expected utility. But arguably even more dangerous are games that have *no*

Mind, Vol. 125 . 499 . July 2016
doi:10.1093/mind/fzv160 Advance Access publication 6 May 2016

expected utility whatsoever.[1] Enter the Pasadena game. Its expectation is a conditionally convergent series, and as such, it can be rearranged to yield any value; hence, it apparently has no expectation at all (Nover and Hájek 2004; Hájek and Nover 2006, 2008; Hájek 2014). Yet it seems that one can still make rational choices involving the game — for example, preferring the Altadena game, which increases each of the Pasadena game's outcomes by a dollar, and thus strictly dominates it. To be sure, these games may be regarded as pathological, residing at the outskirts of the space of decision problems that we might care about (or beyond). But as we will soon see, we can quickly bring them home to infect everyday decision problems that everyone cares about. More generally, options with undefined expectation — expectation *gaps* — rock the very foundations of expected utility theory. It seems that the theory needs more thoroughgoing revision.

Friends of expected utility theory have been quick to rehabilitate it. Fine (2008) shows that, consistent with the axioms of that theory, the Pasadena, Altadena, and St. Petersburg games *can* be valued — although, we would add, consistent valuations may rank them *in that order*, which is surely exactly back to front. Sprenger and Heesen (2011), among others, argue that utilities must be bounded. We disagree (for reasons given in Nover and Hájek 2004); and in any case, as we will see, unbounded utilities are not necessary for expectation gaps. Easwaran (2008) refines the notion of expectation to two notions, according to its role in the weak and strong laws of large numbers, respectively: *weak* and *strong* expectation. The Pasadena and Altadena games have weak expectations (ln 2 and ln 2 + 1, respectively), and arguably should be valued accordingly. This is still very much in the

[1] Broome (1995) judges the St. Petersburg game to be such a game. He writes:

> Part of the reasoning that leads to the St Petersburg paradox is the claim that the expectation of your winnings from the game exceeds any finite amount. But strictly there is no such thing as the expectation of your winnings. Strictly, a distribution with no finite mean does not have a mean that exceeds any finite amount; *it has no mean at all.* (p. 9, our emphasis)

A reason to say this of the St. Petersburg game is that its expectation series diverges. However, it is important *how* it diverges. Its sequence of partial sums $\{S_n\}$ diverges *to* ∞, in the sense that for any $M > 0$, there exists $N \in \mathbb{N}$ such that for all $n > N$, we have $S_n > M$. This is quite different from other ways in which a sequence of partial sums may diverge — for example, with bounded oscillations, or by diverging *to* $-\infty$. The behaviour of the St. Petersburg game's expectation is thus quite different from that of the *Alternating St. Petersburg* game (in which positive and negative payoffs escalate at the St. Petersburg rate), or that of the *Negative St. Petersburg game* (the mirror image of the St. Petersburg game with all payoffs negative). Merely saying that there is no expectation at all for these games misses crucial distinctions among them — *distinctions regarding their expectations.*

spirit of expected utility theory. However, Easwaran himself concedes that there are other games that may not have weak expectations — for example, the Alternating St. Petersburg game, in which positive and negative payoffs alternate.

Bartha (2016) provides an even more telling example of such a game: the *Arroyo game*. Like the Pasadena game, its expectation conditionally converges to ln 2; yet it has no weak expectation. It is more telling, because the two games seem so similar; it is hard to see any reason for taking the Pasadena game seriously without also taking the Arroyo game seriously. The latter game also puts paid to the tempting hypothesis that every game with a conditionally convergent expectation has a weak expectation: we have even a weak expectation gap here. More recently, Easwaran (2014) extends the method of weak expectations to that of 'principal values', yielding a still more general approach in the spirit of expected utility theory. But again, it apparently cannot value the Alternating St. Petersburg game, or judge it to be inferior to a 'sweetened' version of the game in which every payoff is increased by a dollar.

And so we turn to two theories that, in Bartha's words, provide ways of 'making do without expectations': Colyvan's (2008) *relative expectation theory* (RET), and Bartha's *relative utility theory* (RUT). RET is a conservative extension of the finite fragment of expected utility theory: it gives the same results as standard expected utility theory when there are finitely many states and no infinite utilities. RUT too is designed to keep what is good about standard expected utility, then add to that. In this sense, both theories still bear important resemblances to their forebear. We might say analogically (not literally) that these theories *dominate* expected utility theory: in a certain sense, they deliver all the goods of that theory and more besides. In the next three sections, we will set up the 'more besides', rehearse RUT and RET, and explain this 'certain sense' in more detail. We will then compare RUT and RET head-to-head. Does either theory dominate the other (in our analogical sense)? No — we will see that each theory can do some desirable things that the other cannot.

In the subsequent section, we will offer a *nosology* of expectation gaps that are problematic for both theories. There are further diseases out there to which even these descendants of expected utility theory are not immune. The conclusion is a call for more medicine.

2. Expectation gaps and two contagion problems

Starting with an expectation gap, there is an easy recipe for generating more of them. Following Bartha, let's call the propagation of expected value gaps to other decisions *contagion problems*: the disease of a particular gamble's lacking an expected utility easily spreads.

There are two notable kinds of contagion. The first involves *dominance* reasoning. It is exemplified by the Pasadena game and the Altadena game. Starting with the Pasadena game, we sweeten the outcome in every state by \$1 to produce the Altadena game. Since the Pasadena game is an expectation gap, so is the Altadena game. Expected utility theory is silent about the value of each of them. But decision theory ought to say, loudly and clearly, that the Altadena game is better than the Pasadena game, *because it is*! The Altadena game strictly dominates the Pasadena game: in every state of the world, the former yields an outcome strictly better than the latter does. Now, we might add dominance reasoning to expected utility as a separate tool in our decision-theoretic toolbox. (See Colyvan 2006.) However, it would be better to have a unified decision rule that incorporates both expected utility and dominance reasoning. Moreover, decision theory should also say how *much* better the Altadena game is: (the utility of) a dollar. That's the exact amount that one should be prepared to pay to upgrade from the Pasadena game to the Altadena game. But dominance reasoning is also silent about this — it only gives qualitative verdicts, not quantitative.

A second kind of contagion problem is what Bartha calls *the failure of garden-variety decision-theoretic reasoning*, or for short, *garden-variety contagion*. It is exemplified by *the Pizza problem*: choosing between ordering pizza and ordering Chinese food. (See Hájek and Smithson 2012.) This ought to be a straightforward problem, and if decision theory can't handle it, decision theory is in serious trouble. However, suppose that you assign some positive probability, however tiny, to the prospect of playing the Pasadena game after ordering pizza. Then ordering pizza is an expectation gap: it is a gamble, with outcomes *pizza followed by the Pasadena game*, and *pizza not followed by the Pasadena game*. Presumably the probability that you assign to the former outcome is astronomically smaller than the latter. But that does not save the pizza option from contamination. So expected utility theory cannot value ordering pizza, and it cannot even place your pizza ordering on your preference ordering.

Furthermore, Bayesian orthodoxy cannot criticise you for assigning positive probability to your playing the game: you may do so while adhering perfectly to the probability axioms. Indeed, assigning *zero* probability to this contingent event is more liable to ruffle Bayesian feathers. After all, it is unclear why it should be *zeroed* out by your prior,[2] and unclear how it could ever be zeroed out by conditionalization on your evidence. And yet once you assign positive probability to playing the game after having pizza, the pizza option is poisoned decision-theoretically.

This already shows that expected utility theory cannot even represent your choice between pizza and Chinese, still less advise you about it. Indeed, as long as *any* option in a given decision problem of yours is contaminated — however long your list of options may be — you cannot maximise expected utility over all your options. But while we're at it, we might as well contaminate the Chinese option too. For the same reasoning that applies to the pizza option applies to it. We are left comparing a gap with a gap. If you do what decision theory tells you — nothing — your stomach will remain empty.

3. Relative utility theory

The technical details of Bartha's approach are nicely laid out in his two papers (Bartha 2007, 2016). Here we merely sketch the approach, drawing attention to a few key points. RUT starts with preferences rather than a utility function, and its goal is to represent them. The axioms governing preferences in RUT are as for von Neumann and Morgenstern (1944) with two notable exceptions:

(i) The RUT preferences do not need to obey the Archimedean axiom,[3] so RUT can represent (some) preferences involving

[2] Indeed, proponents of *regularity* will insist that it must *not* be. A regular probability function is one that zeroes out only impossible events; anything contingent receives positive probability. Versions of regularity as a constraint on subjective probabilities have been proposed by Kemeny (1955), Shimony (1955, 1970), Jeffreys (1961), Edwards et al. (1963), Carnap (1963), Stalnaker (1970), Lewis (1980), Appiah (1985), Jackson (1987), and others.

[3] The Archimedean axiom says roughly that whenever you have strict preferences among three gambles, there is some weighted average of the best and the worst of these gambles, with *real*-valued weights, that lies at the same place on your preference ordering as the intermediate gamble. More formally, where '$\succ$' is your preference relation, and '$\sim$' denotes indifference for you:

whenever $G3 \succ G2 \succ G1$, there exists a real-valued $0 < \lambda < 1$ such that $[\lambda G3, (1-\lambda)G1] \sim G2$.

infinite utilities,[4] such as Pascal's Wager and the St. Petersburg game.

(ii) The preference ordering relation need not be total. For example, RUT allows for there to be no preference relation between the Pasadena game and the status quo, but it can represent (some) preferences involving the Pasadena game.

RUT employs a generalized arithmetic ratio of utilities and requires three-way comparisons among outcomes. Its basic quantity is the ratio of two differences in utilities when this is defined:

$$\mathcal{U}(A, B; Z) \;=\; (u(A) - u(Z))/(u(B) - u(Z))$$

(although, as we shall see, the official definition is not in terms of a numerical ratio but, rather, preferences between gambles). The 'base point' Z is needed, because there is presumably no such thing as the worst possible option, which could be taken as a natural zero point. There is no utility analogue of zero degrees Kelvin. Nor is the zero point fixed once and for all on a scale that is privileged for some reason or other. There is no utility analogue of zero degrees Celsius, either. When the right-hand side is not defined, the definition of $\mathcal{U}(A, B; Z)$ is given by Bartha's (13) (2016, p. 819).

In a way, the triadic structure of RUT was always implicit in expected utility theory. The zero point and the unit in an expected utility representation are arbitrary; utility values become meaningful only once they have been fixed. Effectively this means that utilities become meaningful only in relation to two choices — we might call them 'B' and 'Z'. The utility of A measures how it subdivides the interval between them. Or going in the other direction, we might interpret $\mathcal{U}(A, B; Z)$ as the utility of A, once we (arbitrarily) assign a utility of 1 to B, and 0 to Z. But as Bartha notes (2016, p. 816), '$\mathcal{U}(A, B; Z) = \alpha$ stands for something like indifference between A and $[\alpha B, (1-\alpha)Z]$, though slightly weaker'. He goes on to explain

This is violated once infinite utilities are introduced. For example, we may suppose that Pascal preferred salvation (with infinite utility) to \$1, and \$1 to the status quo, but that there was no gamble between salvation and the status quo, with real-valued weights, that he regarded indifferently to \$1. Dropping the Archimedean axiom is a way of formally allowing infinite utilities into the picture.

[4] In Bartha's framework, these are interpreted not as absolutely infinite utilities but as preference relations where any gamble offering a positive probability of one outcome is preferred to the other outcome.

what the sense of 'slightly weaker' is: you prefer the compound gamble $[\lambda B, (1 - \lambda)Z]$ to A if $\lambda > \alpha$, and you prefer A to the compound gamble if $\lambda < \alpha$. But it is not required that when $\lambda = \alpha$, you are indifferent between A and the compound gamble.

This triadic structure means that in RUT, we can no longer speak of the expected utility of a single act. Moreover, we can't ask questions about the relationship between two acts, A and B. Suppose A is preferable to B; then we might reasonably ask how much better A is to B. But this does not make sense in Bartha's account (except in a derivative sense). All we can say is that, relative to some third act, Z (the base point), A is preferable to B, but the amount it is preferable is dependent on the choice of Z. More on this later.

Bartha proves a representation theorem: for any set of preferences that obey von Neumann and Morgenstern's axioms without the total ordering and Archimedean axiom, there exists a relative utility function representation. (See Bartha 2016, p. 819 for details.) Any utility function that satisfies certain (weak) ordering conditions, independence, and the compound-gambles condition is also a RUT utility function. But the RUT utility function genuinely extends orthodoxy: there are preferences that the orthodox theory cannot represent, but that RUT can.

So RUT starts with a given agent's preferences, which may violate either the Archimedean axiom (e.g. thanks to the St. Petersburg game) or the total ordering axiom (e.g. thanks to the Pasadena game). As long as the other standard axioms are obeyed, we can represent these preferences in terms of a utility function — the relative utility function — that has some nice properties:

(1) It features in a representation theorem that provides a way of moving between preferences and relative utility.

(2) The resulting theory agrees with ordinary decision theory on standard cases that have a definite value — the ones that can be represented with a one-place utility function.

(3) The theory also represents some preference relations among outcomes that don't have a definite value.

(4) The theory can also deliver the intuitively correct verdicts in some infinite cases that standard decision theory cannot deliver.

4. Relative expectation theory

The core idea in RET is that instead of calculating the expected utilities of actions to determine which action has the greatest expected utility, we calculate the expectation of the differences in utility between two given actions across the various states. (It can be thought of as a kind of expected opportunity cost.)

More formally, we define the relative expected utility of act A_k over A_l, in a decision problem with n states, as:

$$\mathrm{REU}(A_k, A_l) = \sum_{i=1}^{n} p_i(u_{ki} - u_{li})$$

where p_i is the probability associated with state S_i, and u_{ji} are the utilities of the outcomes resulting from act A_j in state S_i. For infinite state spaces we define relative expected utility similarly:

$$\mathrm{REU}(A_k, A_l) = \sum_{i=1}^{\infty} p_i(u_{ki} - u_{li})$$

where the right-hand side absolutely converges,[5] or diverges to infinity or negative infinity. The decision rules are as follows.

RET Decision Rules:

Choose act A_k over act A_l iff $\mathrm{REU}(A_k, A_l) > 0$.

Be indifferent between A_k and A_l iff $\mathrm{REU}(A_k, A_l) = 0$.

A couple of comments about the states are in order. First, the states need to be independent of the acts (by the lights of the agent). If the probability of a given state changes across the acts, there is no way to define the relative expected utility, since the differences in utilities need to be multiplied by the probability of the state in question. This is undefined if there is more than one probability associated with the state. This raises a more general issue about the identification of the states.

For the mathematics of RET to make sense, we must make sense of the same state under different actions. Indeed, an account of what it is for two outcomes to be associated with the same or different states is required, but was not provided in the original paper. This left RET

[5] In Colyvan (2008) it was not stated explicitly that *absolute* convergence was required but in the current context this should be stated explicitly, as Bartha (2016, footnote 12, p. 812) rightly notes.

vulnerable to some problems.[6] For instance, we need to be able to compare bets like the following.

(Bet 1): $5 if a fair coin toss lands heads; nothing otherwise;

(Bet 2): $6 if a fair die toss lands an even number; nothing otherwise.

As things currently stand, RET is silent on this case because there are no states in common across the two bets: 'heads on a toss of a fair coin' and 'even number on a roll of a fair die' are, on the face of it, different states. Yet, we want to say that (Bet 2) is preferable to (Bet 1) by compelling dominance-like reasoning.[7] The obvious move to make here is to stipulate that we can identify 'heads on a toss of a fair coin' and 'even on a roll of a fair die' because they have the same probability, and that's all that matters.

We thus supplement RET with this probabilistic identification of states: the states under one action can be identified with the states under a different action in the same decision problem iff there is a one-to-one correspondence between the two sets of states that maps each state under one action to a state of equal probability under the second action. In short, states of different gambles are identified by virtue of having the same probability assignments, irrespective of the descriptions of the states and irrespective of how the probabilities in question are generated. States are defined by the probability profile.[8] Then we can identify the states across (Bet 1) and (Bet 2). RET then tells us to prefer (Bet 1). And RET, thus construed, also advises indifference between a St. Petersburg game with a fair coin and another St.

[6] Some, but not all, of the problems that Bartha raises for RET turn on this issue as well.

[7] It is interesting to note that dominance reasoning also has this problem: we need to identify the common states across two actions before the rule of dominance can be used. And although the calculation of expected utilities in standard decision theory doesn't require the identification of states across different acts, the standard matrix layout of such decisions does. So, in a sense, the problem of identification of states is something that needs attention anyway, but it is especially important for RET.

[8] Underdetermination will arise from this account of state identity. For example, the state 'heads' of one fair coin toss will be identifiable with the state of either 'heads' or 'tails' of another fair coin toss. But such underdetermination would seem harmless and perhaps even useful.

Our identification method also captures the spirit of Ramsey's (1931) notion of an 'ethically neutral' proposition. The ways in which the outcomes of a gamble may be realised are not valued intrinsically; the agent cares about the states only insofar as she cares about how they probabilistically conduce to the corresponding outcomes. She does not, for example, love seeing an image of the Head of State on a coin while being unmoved by a configuration of dots on a die. So she is indifferent between gambles that substitute 'the coin lands heads' for 'the die lands even', provided the relevant probabilities remain fixed.

Petersburg game with a different fair coin, or with another randomising device that gives the same probability distribution.

In finite cases — with finite utilities and finitely many states — there is no difference between using RET and calculating the expected utilities of acts A_k and A_l, and choosing the act with the greater expected utility. In this sense, RET is a conservative extension of the finite fragment of expected utility theory. The difference comes in cases such as the St. Petersburg game, where there are infinitely many states. Some of these infinite-state cases turn out to be problematic for standard decision theory, yet RET provides intuitively correct answers. For instance, it is straightforward to show that in a choice between the St. Petersburg game and any game that dominates it, RET counsels us to take the dominating game. And returning to the Pasadena game, it is straightforward to show that RET counsels us to take the Altadena game over the Pasadena game and that the difference in expected value is precisely $1. Although the expected utilities of the Pasadena and Altadena games are not defined, the utilities for each of their states are, and so are their differences. Similarly, RET counsels us to take the sweetened Alternating St. Petersburg game, in which every payoff is increased by a dollar, over the Alternating St. Petersburg game. This priority given to dominating acts is no accident: RET is a conservative extension of dominance reasoning.[9]

Colyvan (2008) presented RET as offering different advice from standard decision theory: as advising one, for example, to *prefer* a game that dominates the St. Petersburg game, whereas standard decision theory advises *indifference* between the two (both having infinite expected utility).[10] But it is unclear that standard decision theory offers any such advice. After all, games such as the St. Petersburg game with infinite expected utility violate the Archimedean axiom

[9] To be more precise, it is a conservative extension of dominance understood as follows: act A dominates act B iff in every state the utilities associated with A are never less than the corresponding utilities for B, and in at least one state *with a non-zero probability* the utility of A is greater than the corresponding utility for B. RET agrees with standard decision theory in advising indifference between two otherwise identical gambles but where one gamble has greater utility associated with a probability zero state. Some see this as a problem, and a stronger version of the dominance principle recommends taking the gamble with the greater utility, irrespective of the probability of the state with the greater utility. But we will assume the weaker version in this paper. The difference between the versions will not matter to the cases that we will consider.

[10] It is a common enough view that standard decision theory advises indifference in such cases. Indeed, one of the present authors has suggested as much (Colyvan 2008).

of standard decision theory.[11] In so far as the presence of this game in our preference rankings violates an axiom of standard decision theory, it is unclear that standard decision theory offers any advice at all about whether to prefer the St. Petersburg game or any of its dominating variants. If an axiom is violated, all bets are off as far as decision theory is concerned. Even if we do enrich standard decision theory to allow infinite expected utilities, such as those naturally arising from the St. Petersburg game, there remains an important (mathematical) issue of whether we can treat two infinities as equal in the relevant sense. After all, standard decision theory says to be indifferent between two acts when their expected utilities are *equal*. But when the two expected utilities in question are infinite, it is unclear whether the right thing to say is that the two are equal, or that they are both infinite and no further comparison can be made (since finite arithmetic does not hold here). There are no such complications in RET's treatment of these cases. In light of all this, perhaps a better way of characterising the difference in advice here is that RET gives sensible advice (choose the dominating game), while standard decision theory gives none.

5. Comparing relative utility theory and relative expectation theory

When comparing RUT and RET, it will be useful to look at their similarities and differences, to keep a tally of their formal advantages and disadvantages, and of what they can and cannot do.

5.1 Similarities between RUT and RET

Let's begin with some respects in which RUT and RET are on a par. On neither theory does it make sense to talk of the expected utility of a single act. But RUT and RET recover all the verdicts of expected utility theory in all the straightforward cases. When it comes to decisions that fall outside the scope of expected utility theory, neither RUT nor RET are what Bartha calls 'gap-filling strategies', such as Easwaran's weak expectation approach, his principal values approach, or Fine's method of consistently extending expected utility theory.[12] In particular, RUT

[11] Consider the three prizes \$1, \$2, and playing the St. Petersburg game (where the latter is taken to have infinite value). The continuity axiom requires there to be a probability p such that $p \cdot \$1 + (1\text{-}p) \cdot \$\infty = \$2$. But there is no such p.

[12] Fine notes these possible extensions of expected utility theory, but he does not endorse them as a good way of handling these problems.

and RET cannot compare the Pasadena game and its kin with the status quo, and with options of well-defined expected utility more generally. It is not obvious whether this is a cost or a benefit of these theories; it is not obvious that they *should* be able to make these comparisons. (We agree with Bartha that it is a problem with Fine's method that it allows one to value these games *arbitrarily*, and worse still, we would add, that their values can be *independent* of each other.) This is a cost if the Pasadena game is really worth \$ln 2, as the methods of weak expectations and principal values say it is. (And in that case there is yet another cost of Fine's method: it allows one to *wrongly* value the Pasadena game.) Now, perhaps some future gap-filling method — perhaps *feeble* expectations?! — will correctly tell us what the Alternating St. Petersburg game is worth. In the meantime, the jury is out on how RUT and RET should be judged with respect to these games.

Both RUT and RET do a good job of representing preferences involving options with infinite utilities — up to a point. We have already seen one such option: playing the St. Petersburg game. A different kind of infinite utility is that of salvation, according to Pascal in his famous Wager. We might call this infinite utility *in a single hit*, rather than aggregated, St. Petersburg-style. Indeed, in the St. Petersburg game, the player is certain to get a finite payoff; it is just a quirk of the expectation series, each term of which is finite, that it diverges to infinity. Salvation, by contrast, comes with a guarantee of infinite utility (according to Pascal): it appears as an infinite entry in the Wager's decision matrix of utilities of possible outcomes. Bartha explicitly created RUT to handle Pascal's Wager, and it also yields some happy results regarding the St. Petersburg game S, and its enhanced counterparts S^* and S^{**}, which Bartha considers. (We will see a less happy result in a moment.) Similarly RET does a good job of the problem cases involving infinity that it was designed to handle: for example, prising apart games such as Pasadena and Altadena.

RET and RUT are alike with respect to contagion problems that stem directly from RUT's three axioms. Among other things, they agree that if an option is sweetened in finitely many states, but otherwise left the same, the resulting option should be preferred. For example, Bartha's truncated Altadena games A_n should be preferred to the Pasadena game by both RET's and RUT's lights.

5.2 *Differences between RUT and RET*

Now, let's bring out some differences between the theories. We begin with reasons that apparently favour RET over RUT; then we will look at reasons that do the reverse. We might imagine this as a head-to-head contest of the two theories over various rounds of evaluation, though we will see that their fortunes fluctuate within the rounds, and there will be no decisive winner at the end: we won't see RET rout RUT, or RUT rout RET. We do think that after the dust settles RET is in better shape than RUT overall, but we concede that neither theory is problem-free.

5.3 *RET's advantages over RUT*

RET is simpler than RUT: two argument places are fewer than three. RET, unlike RUT, has no need for a base point Z—whether the expected utility difference between A and B is positive or not depends only on A and B, and it is invariant under positive linear transformations of the utility scale. Happily, though, RUT is invariant under linear transformations: if $\mathcal{U}' = a\mathcal{U} + b$ for constants a and b, then $[\mathcal{U}'(A) - \mathcal{U}'(Z)]/[\mathcal{U}'(B) - \mathcal{U}'(Z)]$ is identical to $[\mathcal{U}(A) - \mathcal{U}(Z)]/[\mathcal{U}(B) - \mathcal{U}(Z)]$. And the flexibility in the choice of Z may have some benefits. Z is like the *z*ero point of a ruler. If A and B are two infinitely good outcomes that we want to compare, we may not want Z to be a merely finitely good outcome—from its perspective, A and B are both infinitely far away, and any distinction between them in value is telescoped. At that distance, so to speak, resolution is lost. It may be better to view them from a comparable vantage point. See, for example, Bartha's discussion of the comparison of two enhanced St Petersburg games S^* and S^{**}, which would be trivialized from a finite standpoint (say, the option of getting a dollar), but which is nicely calibrated from the standpoint of the St. Petersburg game S.

Then again, it may seem odd that RUT's comparison between two options A and B is not an internal relation between them—not one that supervenes just on A and B themselves. To be sure, such a comparison is never *reversed* by changing the base point; (R2) in Bartha's representation theorem assures us of that. (It would not even *be* a representation theorem if $\mathcal{U}$ could *mis*represent a preference between a pair of options.) Nevertheless, to take a particularly striking case, whether A's utility relative to B is 1 or not *can* depend on Z (and Bartha gives an example involving three St. Petersburg-like games in which it does). Seen from one vantage point, A and B appear equally good (in the relative sense); seen from another, they do not.

Similarly, RUT cannot say once and for all, as RET does, that the Altadena game is better than the Pasadena game. It can say this relative to the Bajadena game (the Pasadena game *soured* by a dollar), but not relative to an option whose expected utility is well defined, such as \$1. RUT respects dominance reasoning when only finitely many states are involved; however, it can disrespect such reasoning when infinitely many states are involved, as they are in the Altadena/Pasadena comparison. To be sure, it *can* represent a preference for Altadena over Pasadena. As Bartha writes (2016, p. 822), 'Relative utilities allow us to respect dominance reasoning' here; the glass is half full. But it is also half empty: relative utilities allow us to *dis*respect dominance reasoning. RUT can also represent the *reverse* preference — a *perverse* preference. While we're at it, it can also represent preferences of both of those games over St. Petersburg — even more perverse. RUT is susceptible to the same criticism as expected utility theory is in the face of Fine's result. This is no surprise — after all, RUT is more liberal than expected utility theory, representing all of the preference patterns that the original theory does, and more besides. And since we know from Fine that expected utility theory can represent these perverse preferences, RUT has this result too. Granted, decision theory is supposed to be ecumenical, catering even to those who prefer the destruction of the whole world to the scratching of their finger. But this result is unacceptable. It sanctions self-sabotage — refusing a sure dollar, and much worse. One wants decision theory to be more dictatorial!

Of course, given that RUT cannot say that Altadena is better than Pasadena — period — still less can it say that Altadena is exactly a dollar better, as RET does.[13] RUT has no way of representing the natural judgment about the respective utilities:

$$u(\text{Altadena}) = u(\text{Pasadena}) + u(1), \tag{1}$$

because it does not have a one-place utility function. Nor does it have a two-place utility-difference function. Nor does it even have a version of (1) that replaces the one-place u with three-place relative utilities $\mathcal{U}$. That's because the Pasadena game is not comparable to \$1. Suppose we try:

[13] Thanks to Paul Bartha for help in formulating the rest of this paragraph.

$$\mathcal{U}(\text{Altadena, Pasadena; Bajadena}) = \mathcal{U}(\text{Pasadena, Pasadena; Bajadena}) + \mathcal{U}(\$1, \text{ Pasadena; Bajadena}).$$

The third term is undefined, because \$1 is not comparable to either the Pasadena game or the Bajadena game. And so it goes — no matter which options we pick, no equation like (1) will make sense.

Bartha (2016) shows how to recover the intuitive preference pattern of Altadena over Pasadena and Pasadena over Bajadena (with Pasadena lying exactly in the middle). The trick is to lay down further assumptions, (18)–(20). (19) and (20) are continuity assumptions concerning truncated Bajadena games B_n that 'converge' to the Bajadena game, and truncated Altadena games A_n that 'converge' to the Altadena game. But once we offer counterpart assumptions to Fine, he too can recover our intuitive preference patterns.[14] There's a good sense in which Fine can recover our intuitive judgments to the same extent that Bartha can; hence to the extent that Fine cannot recover them, neither can Bartha.

RUT gets the right result that all of the A_n should be valued above the Pasadena game (thanks to the Independence axiom, which only involves finitely many alterations to a game). However, RUT also allows one to value the Altadena game *below* all the A_n, when intuitively, Altadena should be valued *above* all the A_n. After all, they converge to it 'from below'. The problem, roughly, is that the utility of the limit game (Altadena) can be different from the limit of the sequence of utilities of the A_n. More precisely, we have:

$$\mathcal{U}(\text{Pasadena, Altadena; Bajadena})$$

can be different from:

$$\lim_{n\to\infty} \mathcal{U}(\text{Pasadena}, A_n; \text{ Bajadena}).$$

This is not so obviously a kind of self-sabotage as before, but it is still arguably an unwelcome kind of discontinuity — a discontinuity *at infinity*, as we might say.[15]

[14] We are grateful to Rachael Briggs for this point.

[15] See Bartha, Barker, and Hájek (2013) for further discussion of discontinuities at infinity in decision theory.

RET, on the other hand, doesn't encounter these problems. Although RET does not have a one-place utility function — the RET utility function is two-place — there is a very natural way of rendering (1). In RET we can say that the relative utility of Altadena over Pasadena is \$1: REU(Altadena, Pasadena) = 1. Indeed, this is not just the RET surrogate for (1); it is even better than (1). After all, for all its intuitive appeal, (1) has undefined terms on either side of the equation, so it does not really make sense. The intuitive content of (1), however, is captured by the RET version, which makes perfect sense.

RET can also make good sense of Altadena being better than Pasadena and that being better than Bajadena. Indeed, these orderings are trivial for RET, given that RET preserves dominance reasoning (whenever the utilities for the various states are well defined, as they are in these cases). RET can say that Altadena is better than Pasadena, period. There is no need to appeal to a third game relative to which Pasadena is better than Altadena.

5.4 RUT'S (alleged) advantages over RET

So far so good for RET; but it too has its problems. First, let us note a restriction in RET: as it currently stands, RET assumes the independence of states from acts. In standard decision theory this assumption can be easily dropped by moving to Jeffrey-style decision theory (Jeffrey 1983), but that's not a viable option here. A more serious difficulty with RET is noted by Bartha. RET assumes that the utility of each outcome in the decision matrix is well defined — so it is susceptible to garden-variety contagion problems as in the Pizza case, and to 'single-hit' infinite utilities problems, as we see in Pascal's Wager. RUT scores a clear victory in the latter problems,[16] although they are not as central to our concerns in this paper as the former problems, which we will now explore further.

For instance, Petersen (2011) draws attention to the fact that RET cannot offer advice on a choice between a St. Petersburg-like game but with an infinite payout on one of the outcomes and a similar game with an infinite payout on a higher-probability outcome. Petersen appeals

[16] RUT is tailor-made to handle single hit infinities, at least in comparisons to merely finite outcomes, since infinite utility is not represented as such, but only relationally (the utility of one outcome relative to another, given some base point). Thus RUT is able to avoid the problem that an outcome of infinite utility swamps expected utility calculations, as we find in Pascal's Wager (Bartha 2007). There, RUT can give advice other than 'bet on God', while RET cannot.

here to what we might think of as a better-chances condition — an analogue of the better-prizes condition of von Neumann-Morgenstern utility theory. This better-chances condition encodes an intuitively plausible probability dominance principle, and it advises us that we ought to choose the second game because it gives us a better chance at the infinite prize. RET cannot deliver this advice. Indeed, RET is strangely silent on this and similar cases involving infinities in the decision matrix.

Bartha's example of *Pizza Special* is a choice between pizza and Chinese food but with a small chance of playing the Pasadena game in each case, and with pizza being preferred to Chinese either way. Bartha sets up this example with the same small chance of playing the Pasadena game in each case. This does not give him the result he wants, since the state in which the agent ends up playing the Pasadena game can be decomposed into infinitely many states — the states of the Pasadena game — and each of these will have well-defined utilities.[17] Each outcome will have the utility of the corresponding state of the Pasadena game plus either Chinese food or pizza. The utilities of the Pasadena game will term-by-term disappear (since for each state, the components of the relative utility coming from the Pasadena game will be zero: the difference of two equal terms). This leaves only the utility of pizza and Chinese food, and by construction the former is preferable to the latter. The verdict that RET delivers is the intuitively correct one: choose pizza! Pizza Special is not the counterexample to RET that Bartha wanted.

Bartha is correct, however, that RET assumes well-defined (and finite) utilities for each outcome in the decision matrix. It's just that his choice of example was not the best to exploit this limitation of RET. Bartha needs an example where one cannot get the cancellation across the states of the Pasadena game that we have in Pizza Special. Setting the probabilities of pizza and playing the Pasadena game to be slightly different from Chinese food and playing the Pasadena game will do the trick. So his point about RET being ill equipped to deal with contagion is right; it's just that his main example was one that RET can handle.[18]

According to Bartha (2016), RUT's main comparative advantage over RET is that the former but not the latter can deal with contagion cases. Bartha shows how RUT can deal with the Pizza Special case and deliver

[17] Thanks here to Rachael Briggs and Hanti Lin.

[18] Peterson's (2011) counterexample is along similar lines but he makes sure that the contagious utility — in this case, the expected utility of a St. Petersburg game — is associated with different states and different probabilities so there is no cancellation.

the intuitively correct result of choose pizza. But in his treatment he relies on the fact that the chance of playing the Pasadena game is the same in all three of the actions compared: pizza plus Pasadena, Chinese plus Pasadena, and no meal. As we just saw, however, RET can handle this case. The problem case for RET is where the probability of playing the Pasadena game is different under each action. We are yet to see how RUT can deal with such cases. Indeed all three contagion cases Bartha uses RUT on are also dealt with by RET[19] (and with the same advice in each case). To be sure, RET is susceptible to contagion cases of the kind Bartha alludes to, but at this stage it is not clear whether RUT has a comparative advantage when it comes to these more difficult cases.

That said, there are other problem cases in the vicinity for RET. One particularly intriguing example, due to James Joyce (private communication), puts pressure on our earlier proposal of individuating states via their probabilities. Consider the following decision set up with infinitely many states and three bets under consideration.

In this example we have REU(A, B) = 1 – 1/2 + 1/3 – 1/4 + 1/5 – 1/6

	S1	S2	S3	S4	S5	S6	…
Prob	1/4	1/4	1/8	1/8	1/16	1/16	…
Bet A	\$4	\$0	\$8/3	\$0	\$16/5	\$0	…
Bet B	\$0	\$4/2	\$0	\$8/4	\$0	\$16/6	…
Bet B^*	\$4/2	\$0	\$8/4	\$0	\$16/6	\$0	…

+…, which is only conditionally convergent, so REU(A, B) does not have a value. Now permute the outcomes in B among equally probable states to get bet B^*. REU(B^*, B) = ½ - ½ + ¼ - ¼ + 1/6 – 1/6 +…. This too is only conditionally convergent, so it does not have a value either. But since we are stipulating that states are to be individuated by their probabilities, B and B^* should be indistinguishable and REU(B, B^*) should be zero. Equally puzzling is the fact that REU(A, B^*) *does* have a determinate value: REU(A, B^*) = ½ + 0 + 1/12 + 0 + 1/30 + 0 + 1/56 + 0 +…, which is absolutely convergent. It is very strange that REU(A, B^*) has a determinate value while REU(A, B) does not. After all, if states are individuated only by their probabilities, it looks as though we should be able to compare A and B. Here are two ways:

[19] At least RET supplemented in the way we suggested earlier by identifying states via probability profiles can deal with such problems. In fairness to Bartha, though, this account of state identification was not in the original presentation of RET, and that presentation is his target.

comparing the first state in A with the first state in B, the second state in A with the second state in B, the third state in A with the third state in B, the fourth state in A with the fourth state in B,...

comparing the first state in A with the second state in B, the second state in A with the first state in B, the third state in A with the fourth state in B, the fourth state in A with the third state in B,...

But, the first comparison is REU(A, B) and the second is REU(A, B^*). So, it looks as though either the states have to be individuated by something more than their probabilities, or we have to swallow the above paradoxical results about RET's verdicts about the comparisons between A, B and B^*.[20]

This is a serious problem for RET, and it arises primarily because of RET's identification of states via probability profiles. Recall that RET needed this account of state identification to deliver various compelling dominance-like verdicts. But perhaps we've jumped from the frying pan into the fire here. Be that as it may, the identification of states via probabilities is very natural and we suspect is in the background in many applications of dominance-like reasoning.[21] It is worth seeing how far we can proceed with this account of state identification, perhaps supplementing it with something a little more sophisticated.[22] Clearly more work is needed in order for RET to circumvent problem cases such as Joyce's.

For present purposes, however, what is important is whether RUT has a comparative advantage here. RUT is potentially in better shape with regard to Joyce-style variants of the Pasadena paradox because RUT does not rely on probability profiles to identify states. As we've

[20] It is worth noting that similar problems relating to the individuation of the states arise for dominance reasoning. For example, A dominates B^* but not B.

[21] Recalling an earlier example, we reason this way when we determine that (Bet 2) is preferable to (Bet 1):

(Bet 1): \$5 if a fair coin toss lands heads; nothing otherwise;

(Bet 2): \$6 if a fair die toss lands an even number; nothing otherwise.

The idea that it's the probability profiles over outcomes that matter, rather than exactly how they are realized, underpins *stochastic dominance*, a generalisation of dominance. Gamble A stochastically dominates gamble B if for any outcome x, A gives at least as high a probability of paying out at least x as B does, and for some x, A gives a higher probability of paying out at least x.

[22] Or perhaps the option of biting the bullet has more plausibility than first appears. We might, for example, accept that although B and B* are equivalent gambles, we need not be indifferent between them. Seidenfeld, Schervish, and Kadane (2009) prove some interesting results in this direction. They see this as one of the prices you pay for violating the Archimedean axiom.

already stressed, this identification is at the heart of the problem here. If RUT can handle such cases, that would count as a significant point in its favour. We leave it as a challenge to Bartha (or any other defenders of RUT) to show that RUT can deal with such cases. And the challenge for friends of RET is to give a more sophisticated account of identification of states — one that does not succumb to such problem cases.

Stepping back for a moment, the different costs and benefits of RET compared with RUT can't be fully explained by the one taking ratios while the other takes differences, or even by the one being a function of three options while the other is a function of two. Crucial is the difference in their axiomatic bases. While still dropping total ordering and the Archimedean axiom, Bartha could add an axiom that supports dominance reasoning in full generality — even over infinitely many states.[23] This enriched theory may well deliver a similar representation theorem to his own.

6. A nosology of decision-theoretic diseases

So far we have shown that neither of the two decision theories under consideration, RUT and RET, dominates the other, in the sense that neither can do all the other does and more besides. Each has its strengths and weaknesses. This prompts the search for a still broader theory, one that dominates *both* these theories — sharing their strengths, while avoiding their weaknesses. Now, if the Pasadena game and its kin were the only ills that threatened decision theory, we might be tempted to endorse both RET and RUT, depending on the particular decision problem at hand, and be content with that. But, alas, things are much worse once the full range of potential diseases is appreciated. A fully adequate theory has to do *new* work that neither of these theories can do.

Expected utilities are sums of products of utilities and probabilities. If the *utilities* and *probabilities* are *well-defined real numbers*, ordinary real-number arithmetic applies. However, if any of these conditions fail, we have a recipe for problems for expected utility theory, in which utilities are numerical quantities.

So how could the conditions fail? The contrast to a quantity's being well-defined is its being *undefined* — and this could be the case for either a utility or a probability. The most interesting and relevant contrasts to a quantity's being a real number are its being

[23] Thanks here to Paul Bartha (personal communication).

(positive or negative) *infinite* or *infinitesimal*—the former could be the case for a utility, the latter for a probability. Of course, there cannot be infinite probabilities. There's a sense in which a utility can be infinitesimal, but this will not be invariant under positive affine transformations (much as its being 0 will not be), so we will set this aside. However, a utility's being (positive or negative) infinite *is* invariant, and those are the cases we will focus on. There is a growing literature (e.g. Bartha and Hitchcock 1999; Hájek 2003; Elga 2004; Wenmackers 2013) on infinitesimal probabilities, though not much discussion of their role in expected utilities. Where an infinitesimal (as opposed to 0) probability makes only an infinitesimal difference to an expected utility, it might be dismissed as a 'don't care'. The interesting cases are ones in which it makes a finite, or even infinite difference. Furthermore, problem cases can also be achieved by the *combination* of utilities and probabilities, even when they are individually unproblematic.

Given this taxonomy of problem cases, examples virtually write themselves:

1 *Utility*

1.1 Undefined utility

As we have seen, a utility gap that is plugged into any gamble makes that gamble's expectation gappy. Bartha's *Pizza Special* is such a case.[24] Here the gap arises from the undefined expectation of the Pasadena game. But we might also entertain more direct sources of undefined utilities. For example, if there are incommensurable values, then they give rise to utility gaps. Take two outcomes O_1 and O_2 that are incommensurable with respect to value. If O_1 is a utility gap, we are done. If it is not, then O_2 must be. For otherwise it would have a utility, one that it is either greater than, less than, or equal to the utility of O_1;[25] and thus O_1 and O_2 would be commensurable after all. Perhaps, for instance, friendship is

[24] This is the choice between Pizza with a small chance of playing the Pasadena game and Chinese food with the same small chance of playing the Pasadena game, where Pizza is preferable when the choice is restricted to those states where the Pasadena game is not played, *and* when restricted to the state where it is.

[25] We may set aside here exotic utility theories in which utilities are not numbers—for example, theories in which they are vectors. Such theories have already abandoned expected utility theory in the sense that we intend.

incommensurable with monetary outcomes, in which case they cannot all be put on a numerical utility scale.

1.2 Infinite utility (positive and negative)
Infinite utility can arise in two ways:

1.2.1 In a 'single hit'
Think again of Pascalian salvation, in which a single outcome has infinite utility. Negative infinite utility can also arise this way: damnation, regarded as a single outcome with negative infinite utility. More generally, a 'single-hit' infinite utility is an infinite value (positive or negative) of an outcome in a decision matrix, as opposed to an infinite expectation of an option with only finite values attached to the outcomes, as in the St. Petersburg game. Such 'single-hit' infinities give rise to a variety of problems. Perhaps the most serious is that they swamp any non-zero probability, in the sense that a small chance at salvation has the same expected utility as guaranteed salvation. This is not merely a theoretical problem either. Some authors take (some features of) the natural environment to be infinitely valuable, and thereby cannot say, with standard decision theory, that a strategy with a greater chance of saving the natural environment is preferable to a strategy with a smaller chance of success (Colyvan, Justus, and Regan 2010).[26]

1.2.2 As the result of combining finite utilities
Here the poster child is the St. Petersburg game. Negative infinite utility can also arise this way: the *negative* St. Petersburg game, in which escalating gains are replaced by escalating losses. (Or simply consider switching places with the bookie offering the St. Petersburg gamble.) These infinite expectations are problematic for reasons similar to the single-hit infinities. Apart from the familiar point that an infinite value is too high for a game that guarantees only a finite payout, the infinite expectation makes it impossible in standard theory to discriminate between different games with infinite expectations. We've already seen cases of this: St. Petersburg versus the enhanced games S^* and S^{**}.

[26] Such problems also arise when trying to make sense of Kantian duties and prohibitions in the context of combining a deontological ethical theory with decision theory (Colyvan, Cox, and Steele 2010).

2 *Probability*

2.1 Undefined probability

Here is the route to expectation gaps that we want to emphasise — cases in which the relevant *probabilities* are gappy. Probability gaps famously arise as so-called 'non-measurable sets'. In some probability spaces, certain symmetry constraints are incompatible with all events having probabilities. This follows from counterpart results in measure theory — for example, Vitali's proof of the existence of non-measurable subsets of the [0, 1] interval (Halmos 1950), or the Banach-Tarski paradox (Wagon 1985). As a vivid example, consider the random selection of a point from the [0, 1] interval. The selection is unbiased; this corresponds to the translation-invariant probability distribution, Lebesgue measure. Famously, certain subsets do not receive any probability at all; call one of them *N*. Now, imagine that you are offered a bet that will pay a million dollars if the selected point lies inside *N*, nothing otherwise. How much should you pay for the bet? Expected utility theory cannot tell you, as the expected utility of this bet is undefined.[27]

The existence of non-measurable sets can be rigorously proven, given certain assumptions (notably countable additivity and the axiom of choice); but such events are difficult to countenance. Some putative examples of probability gaps are easier to countenance, but also more controversial. Perhaps free choices are *objective chance* gaps, in virtue of the metaphysics of freedom. And as such, they may be reasonable candidates for subjective probability gaps for principled reasons — indeed, for Principal Principle reasons.[28] According to Bayesian orthodoxy, undefined probabilities

[27] Perhaps we can at least put bounds on the expected utility. The probability of N is bounded below by its inner measure m_i, and bounded above by its outer measure m_O. Then plausibly the expected utility is bounded below by $u(1{,}000{,}000)m_i$, and bounded above by $u(1{,}000{,}000)m_O$.

[28] According to the Principal Principle (suppressing some complications), a rational agent's credences should reflect corresponding chances:

$P(X \mid \text{chance}(X) = p) = p.$

The Principle says nothing about chance *gaps*, but an analogous principle is somewhat plausible:

$P(X \mid \text{chance}(X)$ is undefined) is undefined. (See Hájek 2003 and MS.)

may arise by conditionalizing on a probability zero event. For example, according to standard probability (and measure) theory, the probability of a randomly selected point on the Earth's surface being in the Eastern hemisphere, given that it is on the equator, is undefined (rather than the intuitively reasonable answer of $1/2$).[29] Finally, probability gaps may be completely mundane: for whatever reason, you simply do not assign a probability to *X*. Then *X* is an expectation gap for you.

Note that bounding the utility function will not help here. There need be nothing problematic about the utilities in such cases — the utility of a million dollars, or what have you. So one popular response to the St. Petersburg game, the Pasadena game, and their ilk, will not solve these problems.

2.2 Infinitesimal probability

Events of infinitesimal probability will typically make infinitesimal contributions to expected utilities, and so they might be dismissed for this reason. However, when they combine with (positive or negative) infinite utilities, problems may arise. In Pascal's Wager, suppose that you assign infinitesimal probability to God's existence; what is your expected utility for wagering for God? That depends on how the size of the infinitesimal probability compares to the size of the infinite utility of salvation (Hájek 2003). The answer could be infinite, finite, or infinitesimal.

3 *Utility and probability in combination*

Let us return to where we began: the Pasadena game. All of the utilities and probabilities are unproblematic on their own: they are well-defined real numbers. It is their combination that is problematic. Indeed, it is this feature of the Pasadena game, and the St. Petersburg game for that matter, which makes it particularly disturbing. No heavy-handed tactics such as banning completed infinities or incommensurability will work; all that is required for problems to occur is that the utility function be unbounded. All the utilities associated

[29] This might be thought of as a shortcoming of the standard definition of conditional probability and can be avoided if a different definition of conditional probability is adopted, such as one that takes conditional probability to be primitive (Hájek 2003; Popper 1959; Roeper and Leblanc 1999). Be that as it may, this is still a problem for standard decision theory using classical probability theory.

> with the outcomes in both the St. Petersburg game and the Pasadena game are finite, commensurable, and well-defined; it's the resulting expectation that's problematic. This is not a case of garbage in, garbage out.

And so it goes with all of the contagion problems. Given our nosology, we now see a general recipe for contaminating expected utilities. Start with a given expectation gap: this will be our pathogen. Then either plug it into a new compound gamble, or alter all its utilities for outcomes by a fixed amount; the result will be another gap. For example, start with the bet on whether the randomly selected point lies in *N*. Now sweeten the bet — say, let the prize be a million dollars *plus one*. Or suppose that you win the prize if the point lies in some *superset* of *N*. Either way, dominance reasoning tells you that you should prefer the sweetened bet. But expected utility theory stays silent.

Despite Bartha's important contribution, there is still much ado for decision theory when expectations are undefined. And, indeed, there are problems that go beyond expectation gaps. Neither RUT nor RET, as they stand, can deal with all the problems above. To be sure, both RUT and RET make some progress. And it is interesting that they make progress in different ways, but both fall short of what we want.

7. You can't always get what you want

It is fair to say that no single decision theory currently on the table is able to deal with all known problem cases. We may want such a theory, but as the Rolling Stones advised us back in 1969, you can't always get what you want. But perhaps we can get what we need: some decision theory or other that deals with each problem case. We hoped for a universal elixir — a decision theory that would cure all known diseases in rational choice theory. Instead, we may have to settle for case-by-case treatments of the diseases in question. There are no universal elixirs in medicine, so why would we expect them in decision theory?

If we at least had *some* cure for each disease — a decision tool fit for each decision-theoretic pathology — we could use the suite of decision tools to get by, albeit in a disunified, ad hoc, and thus rather unsatisfying way. This raises a further problem. If we don't have a principled method for delineating the applications of the different decision theories, how do we know that we are using the right decision rule? Absent a unified theory, we require a decision rule to determine which decision rule to use in a given situation. We would need, as it were, a decision rule for decision rules. But at this stage we don't even have decision theories to cover each of the problem cases, so we needn't worry about choosing between the decision rules yet. First we need solutions to all the problem cases in §6.

Bartha's RUT is a very useful and interesting new tool to add to the decision theory tool kit. It doesn't do it all, and Bartha never claimed that it did. But it is an important step towards getting what we need, if not what we want.[30]

References

Appiah, Anthony 1985: *Assertion and Conditionals*. Cambridge University Press, Cambridge.

Arnauld, Antoine and Pierre Nicole 1964: *The Art of Thinking; Port-Royal Logic*. Translated, with an introduction by J. Dickoff and P. James, and a foreword by C. W. Hendel, Bobbs-Merrill, Indianapolis (first published in 1662).

Bartha, Paul 2007: 'Taking Stock of Infinite Value: Pascal's Wager and Relative Utilities'. *Synthese*, 154(1), pp. 5–29.

—— 2016: 'Making Do Without Expectations'. *Mind*, 125(499), pp. 799–827.

Bartha, Paul, John Barker, and Alan Hájek 2013: 'Saint Peter, Saint Petersburg, and Satan'. *Synthese*, December, DOI 10.1007/s11229-013-0379-9.

Bartha, Paul, and Christopher Hitchcock 1999: 'The Shooting-room Paradox and Conditionalizing on Measurably Challenged Sets'. *Synthese*, 118(3), pp. 403–37.

Broome, John 1995: 'The Two-envelope Paradox'. *Analysis*, 55(1), pp. 6–11.

[30] Thanks to Rachael Briggs, David Chalmers, Seth Lazar, and Hanti Lin for helpful discussions on the topic of this paper and to Edward Elliott, James Joyce, and Ralph Miles for comments on earlier drafts. Special thanks to Paul Bartha, who was a constant source of instruction and insight in the course of our writing this paper — at least in this case, we got all that we could have wanted. This work was funded by the Australian Research Council: a Discovery Grant to Alan Hájek (grant number DP130104665) and a Future Fellowship grant to Mark Colyvan (grant number FT110100909).

Carnap, Rudolf 1963: 'Carnap's Intellectual Autobiography'. In P. A. Schilpp (ed.), *The Philosophy of Rudolf Carnap*, The Library of Living Philosophers, Vol. XI, Open Court, La Salle, pp. 1–84.
Colyvan, Mark 2006: 'No Expectations'. *Mind*, 115(459), pp. 695–702.
—— 2008: 'Relative Expectation Theory'. *Journal of Philosophy*, 105(1), pp. 37–44.
Colyvan, Mark, Damian Cox, and Katie Steele 2010: 'Modelling the Moral Dimension of Decisions'. *Noûs*, 44(3), pp. 503–29.
Colyvan, Mark, James Justus, and Helen M. Regan 2010: 'The Natural Environment is Valuable but not Infinitely Valuable'. *Conservation Letters*, 3(4), pp. 224–8.
Easwaran, Kenny 2008: 'Strong and Weak Expectations'. *Mind*, 117(467), pp. 633–41.
—— 2014: 'Principal Values and Weak Expectations'. *Mind*, 123(490), pp. 503–16.
Edwards, Ward, Harold Lindman, and Leonard J. Savage, 1963: 'Bayesian Statistical Inference for Psychological Research'. *Psychological Review* 70(3), pp. 193–242.
Elga, Adam 2004: 'Infinitesimal Chances and the Laws of Nature'. *Australasian Journal of Philosophy*, 82(1), pp. 67–76.
Fine, Terrence L. 2008: 'Evaluating the Pasadena, Altadena, and St. Petersburg Gambles'. *Mind*, 117(467), pp. 613–22.
Gibbard, Allan and William Harper 1981: 'Counterfactuals and Two Kinds of Expected Utility'. In W. Harper, R. Stalnaker, and G. Pearce (eds.), *Ifs: Conditionals, Belief, Decision, Chance, and Time*, Reidel, Dordrecht, pp. 153–90.
Hájek, Alan 2003: 'What Conditional Probability Could Not Be'. *Synthese*, 137(3), pp. 273–323.
—— 2006: 'In Memory of Richard Jeffrey: Some Reminiscences and Some Reflections on The Logic of Decision'. *Philosophy of Science*, 73(5), pp. 947–58.
—— 2012: 'Pascal's Wager'. In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2012 Edition), URL = < http://plato.stanford.edu/archives/win2012/entries/pascal-wager/ > (first published in 1998).
—— 2014: 'Unexpected Expectations'. *Mind*, 123(490), pp. 533–567.
—— MS: 'Staying Regular?'
Hájek, Alan and Harris Nover 2006: 'Perplexing Expectations'. *Mind*, 115(459), pp. 703–20.
—— 2008: 'Complex Expectations'. *Mind*, 117(467), pp. 643–64.

Hájek, Alan and Michael Smithson 2012: 'Rationality and Indeterminate Probabilities'. *Synthese*, 187(1), pp. 33–48.
Halmos, Paul R. 1950: *Measure Theory*. Springer, New York.
Jackson, Frank 1987: *Conditionals*. Blackwell, Oxford.
Jeffrey, Richard C. 1983: *The Logic of Decision*, 2nd edition. University of Chicago Press, Chicago (first edition: 1966).
Jeffreys, Harold 1961: *Theory of Probability*, 3rd edition. Clarendon Press, Oxford.
Joyce, James 1999: *The Foundations of Causal Decision Theory*. Cambridge University Press, New York.
Kemeny, John G. 1955: 'Fair Bets and Inductive Probabilities'. *The Journal of Symbolic Logic*, 20(3), pp. 263–73.
Lewis, David 1980: 'A Subjectivist's Guide to Objective Chance'. In R. C. Jeffrey (ed.), *Studies in Inductive Logic and Probability II*, University of California Press, Berkeley, pp. 263–93.
—— 1981: 'Causal Decision Theory'. *Australasian Journal of Philosophy*, 59(1), pp. 5–30.
Nover, Harris and Alan Hájek 2004: 'Vexing Expectations'. *Mind*, 113(450), pp. 237–49.
Nozick, Robert 1969: 'Newcomb's Problem and Two Principles of Choice'. In N. Rescher (ed.), *Essays in Honour of Carl G. Hempel*, Reidel, Dordrecht, pp. 114–46.
Peterson, Martin 2011: 'A New Twist to the St. Petersburg Paradox'. *Journal of Philosophy*, 108(12), pp. 697–99.
Popper, Karl 1959: *The Logic of Scientific Discovery*. Hutchinson & Co, London.
Ramsey, Frank P. 1931: 'Truth and Probability'. In R. B. Braithwaite (ed.), *The Foundations of Mathematics and other Logical Essays*, Routledge and Kegan Paul, London (first published in 1926), pp. 156–84.
Roeper, Peter and Hugues Leblanc 1999: *Probability Theory and Probability Logic*. University of Toronto Press, Toronto.
Savage, Leonard J. 1954: *The Foundations of Statistics*. New York, Wiley.
Seidenfeld, Teddy, Mark J. Schervish, and Joseph B. Kadane 2009: 'Preference for Equivalent Random Variables: A Price for Unbounded Utilities'. *Journal of Mathematical Economics*, 45(5–6), pp. 329–40.
Shimony, Abner 1955: 'Coherence and the Axioms of Confirmation'. *Journal of Symbolic Logic*, 20(1), pp. 1–28.

——1970: 'Scientific Inference'. In R. G. Colodny (ed.), *The Nature and Function of Scientific Theories*, University of Pittsburgh Press, Pittsburgh, pp. 79–172.

Sprenger, Jan and Remco Heesen 2011: 'The Bounded Strength of Weak Expectations'. *Mind*, 120(479), pp. 819–32.

Stalnaker, Robert C. 1970: 'Probability and Conditionals'. *Philosophy of Science* 37(1), pp. 64–80.

von Neumann, John and Oskar Morgenstern 1944: *Theory of Games and Economic Behavior*. Princeton University Press, Princeton.

Wagon, Stan 1985: *The Banach-Tarski Paradox*. Cambridge University Press, Cambridge.

Wenmackers, Sylvia 2013: 'Ultralarge Lotteries: Analyzing the Lottery Paradox Using Non-standard Analysis'. *Journal of Applied Logic*, 11(4), pp. 452–67.

