

Infinite aggregation and risk*

Abstract

For aggregative theories of moral value, it is a challenge to rank worlds that each contain finitely many valuable events. And, although there are several existing proposals for doing so, few provide a cardinal measure of each world's value. This raises the even greater challenge of ranking *lotteries* over such worlds—without a cardinal value for each world, we cannot apply expected value theory. How then can we compare such lotteries? To date, we have just one method for doing so (proposed separately by Arntzenius, Bostrom, and Meacham), which is to compare the prospects for value at each individual location, and to then represent and compare lotteries by their expected values at each of those locations. But, as I show here, this approach violates several key principles of decision theory and generates some implausible verdicts. I propose an alternative—one which delivers plausible rankings of lotteries, which is implied by a plausible collection of axioms, and which can be applied alongside almost any ranking of infinite worlds.

Keywords: *infinite value; infinite utility streams; decision theory; stochastic dominance; aggregation; Pareto;*

Word count: 8,566

1 Introduction

Consider *aggregative* theories of moral value: those that say that one outcome is at least as good as another if and only if it has at least as great a total aggregate of value, impartially construed.¹ These theories of value include those used by total utilitarianism, total prioritarianism (see Parfit, 1997), and critical level theories (e.g. Blackorby et al., 1995).

*[Acknowledgements omitted for blind review]

¹This definition excludes narrow person-affecting views, as well as any theory under which value does not admit an additively separable representation (e.g., egalitarianism, maximin, averaging). This exclusion is not because such views escape the infinitarian worries described below—typically, they don't—but simply for brevity.

One powerful objection against such theories is this: in a physically realistic setting, they seem to deliver absurd verdicts. To see why, note that our current understanding of cosmology suggests that our universe is *infinite*—it will contain an infinite volume of space and time, as well as infinitely many tokens of every physically possible small-scale phenomenon.² But some physical phenomena are morally valuable (or disvaluable), e.g., perhaps a human brain experiencing pleasure (or pain) for some duration. Our universe will contain infinitely many such events, and so infinitely many instances of value (of value at least some $\epsilon > 0$ on whichever cardinal scale we use).³ If we sum the value of such events in a given outcome, the total will be infinite. And if we include events of negative value too, the total will be *undefined*. So, since our current understanding of cosmology implies that there will occur infinitely many of every such event, valuable or disvaluable, the total sum of value will always be undefined. And one undefined sum is no greater than another. This is bad news for aggregative theories—in practice, they seem unable to say that any available outcome is better than any other. But this is absurd, so perhaps we must reject aggregative theories entirely.

This might be too hasty. Instead of rejecting all aggregative theories, perhaps we can adopt one that avoids this problem of widespread incomparability. By tweaking our method of aggregation, and comparing worlds based on something other than a single real-valued total *sum*, perhaps we can avoid the problem. Fortunately, we have various proposals for doing this. We have the basic proposal of *additivism* from Lauwers and Vallentyne (2004, p. 21), described below. We have the *expansionist* methods of Vallentyne and Kagan (1997, p. 17), Arntzenius (2014, p. 56), and Wilkinson (2021). We have Bostrom’s 2011, pp. 27-30 proposal of using hyperreal numbers. And we have Jonsson & Voorneveld’s 2018 *limit-discounting* method, among various others. I’ll remain mostly agnostic here on which, if any, of these proposals is correct.

Whichever proposal we adopt, we face further problems. As mere human agents, we are uncertain of the outcomes our choices produce, for all choices we ever make. To make decisions in practice, comparing infinite outcomes is not enough; we need to compare *lotteries* over such outcomes.

²This is implied by the widely accepted *flat-lambda* model of cosmology (see Wald, 1983; de Simone et al., 2010; Carroll, 2017, for discussion). It is also implied by the *inflationary view* (see Guth, 2007; Garriga and Vilenkin, 2001). But, by the latter theory, the universe as a whole may have infinite volume but only a finite volume of it within our causal future. If so, it may be physically impossible to cause changes in value at infinitely many different locations, and so the problems raised below may not arise. (But they may still arise if we adopt an *evidential* decision theory for moral decision-making—see MacAskill et al. (n.d.).)

³Specifically, we will have a *countably* infinite number of them. Why? Because they are each positioned in a four-dimensional spacetime. They’ll also each occupy some (exclusive) finite region of spacetime—e.g., for a human brain to experience some quantity of pleasure, it requires some non-zero spatial volume and some non-zero, finite duration. So we can only fit a countably infinite number of those token events into the world.

The standard approach for converting moral comparisons of outcomes into comparisons of lotteries, at least for finitely-valued outcomes, is *expected value theory*: one lottery is at least as (instrumentally) good as another if and only if the *expectation* of its total value is at least as great (see, e.g., Jackson, 1991).⁴ This works fine if the total value of every outcome is finite and so can be represented on a cardinal scale. But it often won't work when outcomes contain infinitely many instances of value—we cannot simply take the expectation of undefined or infinite total values (at least not while producing plausible verdicts). And none of the proposals listed above resolve this—they don't provide total aggregates that can be represented cardinally; they merely give conditions for an outcome to be better than another.

How then can we compare lotteries over outcomes with infinite or undefined value? On existing views, we cannot simply apply expected value theory. But one alternative approach has been proposed, by Bostrom (2011, pp. 27-30), Arntzenius (2014, pp. 53-7), and Meacham (2020), which I describe in Section 3 below. Unfortunately, this solution fails to satisfy several crucial desiderata for any plausible ranking of lotteries, as I demonstrate in Section 4. So a new approach is needed.

I offer such an approach in Section 5: the *expectations of differences* approach. Under this approach, we can adhere to the spirit of expected value reasoning without explicitly assigning expected values to lotteries and comparing them (the traditional, 'differences of expectations' approach). My approach does so while still satisfying key desiderata (unlike the rival Arntzenius-Bostrom-Meacham proposal), and is implied by the conjunction of some highly plausible axioms. And perhaps its greatest advantage is that it can be applied alongside almost any aggregative method of comparing outcomes—one could adopt any of the methods listed above and still use my approach to extend one's comparisons from outcomes to lotteries.

2 Comparing worlds

The challenge is to extend aggregative theories to compare lotteries over outcomes containing infinitely many instances of value. But, to start with, how should we represent such outcomes? And what basic conditions must aggregative theories satisfy when comparing them?

⁴Note that *expected value* is distinct from the frequently-used notion of *expected utility*. And *expected value theory* is distinct from *expected utility theory*, another widely accepted view in normative decision theory. Under *expected utility theory*, *utility* is given by some (indeed, *any*) increasing function of value—perhaps a concave function, such that additional value contributes less and less additional utility. Maximising the expectation of utility may then often give verdicts different from *expected value theory*.

Aggregative theories compare outcomes based on their total aggregates of value. A total aggregate here is some impartial combination of all of the instances of value an outcome contains—its *local values*. But how do we identify and demarcate local values? Following others in the literature, I assume that each local value is fundamentally associated with a *location*: a token entity of some common type that can exist (or have unique counterparts) across different worlds. These locations might be persons, or person-time-slices, or positions in space and time, or something else. They might have some essential and natural topological structure (as spacetime positions do), or they might not (as persons do not). There is disagreement in the literature on the correct such type, and whether they have such structure.⁵ Here I remain agnostic on those questions.

Whatever the relevant type of locations, each outcome will contain some plurality of them, given by the set $\mathcal{L} = \{l_a, l_b, l_c, \dots\}$. And the value at each location is given on a cardinal scale by a function $V : \mathcal{L} \rightarrow \mathbb{R}$. Each outcome (or *world*) can then be represented by an ordered pair $\langle \mathcal{L}, V \rangle$. I'll often use the subscript of each W_i to identify it with its set of locations \mathcal{L}_i and value function V_i . And where all outcomes available in a decision share the same locations, I denote the common location set as \mathcal{L} .

To compare worlds, we need an 'at least as good as' relation. This will be a binary relation \succsim on the set \mathcal{W} of (epistemically) possible worlds. I assume that this relation is both reflexive and transitive⁶ over \mathcal{W} . (Equivalently, we could say that \succsim is a *preorder* on \mathcal{W} .) Strict betterness is denoted by the asymmetric component \succ , and equality by the symmetric component \sim .

But when is one world at least as good as another; when does $W_a \succsim W_b$ hold? This depends on exactly which of the many proposed \succsim relations we endorse. But all of the proposals I mentioned above agree in some cases, including the following.

Suppose you can either rescue one person from painful death or rescue five others. Suppose that your choice has no other effects, and that exactly the same locations (persons, spacetime positions,

⁵For a defence of adopting persons as the appropriate type, see Askill (2019). For arguments in favour of adopting spacetime positions, see Wilkinson (2021) and Author (n.d.).

⁶Transitivity of moral betterness has its critics, e.g., Temkin (2014). It has also received compelling defences from, e.g., Broome (2004); Huemer (2008); Nebel (2018); Dreier (2019). I find it overwhelmingly plausible so, in keeping with the infinite aggregation literature to date, I will assume without argument that it holds.

or what have you) exist either way. We can represent the two outcomes with $W_{\text{save } 1}$ and $W_{\text{save } 5}$.

	l_a	l_b	l_c	l_d	l_e	l_f	l_g	l_h	l_i	l_j	\dots
$W_{\text{save } 1}$:	1	0	0	0	0	0	1	1	1	1	\dots
$W_{\text{save } 5}$:	0	1	1	1	1	1	1	1	1	1	\dots

We cannot compare the total sums of value of these worlds. But we can compare their values at each location, and also at each finite set of locations. For instance, take the first six locations (l_a to l_f); $W_{\text{save } 1}$'s subtotal is 1, and $W_{\text{save } 5}$'s is 5. At all other locations, the worlds have equal value, so let us ignore those. Where they do differ, W_5 has the greater subtotal of value. Thus we might justify $W_{\text{save } 5}$ as the better world.

This is how *Additivity* asks us to reason. Put precisely:

*Additivity*⁷: For any worlds W_a and W_b with the same locations \mathcal{L} , $W_a \succcurlyeq W_b$ if

$$\sum_{l \in \mathcal{L}} (V_a(l) - V_b(l)) \geq 0$$

(either by converging unconditionally, or by diverging unconditionally to $+\infty$).

Note that Additivity is not biconditional—some worlds may be better than others even though Additivity doesn't hold. After all, Additivity by itself does not give us a complete betterness \succcurlyeq relation. Nor does it give a \succcurlyeq relation that compares even a significant fraction of the pairs of worlds our actions might produce in practice. It is just a minimal condition, and one that is satisfied by almost every proposal in the literature.

It is also well-justified—a \succcurlyeq relation that violated it would be a strange betterness relation indeed. As (Lauwers and Vallentyne, 2004, p. 39) show, any transitive \succcurlyeq which satisfies all three of the following highly plausible principles will also satisfy Additivity.

Pareto Over Locations: For any worlds W_a and W_b with the same locations \mathcal{L} , if for *all* $l \in \mathcal{L}$, $V_a(l) \geq V_b(l)$, then $W_a \succcurlyeq W_b$.

If, as well, some $l_i \in \mathcal{L}$ has $V_a(l_i) > V_b(l_i)$, then $W_a \succ W_b$.

⁷This principle is presented and defended by Vallentyne and Kagan (1997, p. 11), Lauwers and Vallentyne (2004, p. 21), and Basu and Mitra (2007).

This says that, if one world has contains at least as much value as another at *every single* location (whatever the relevant type of locations is), then it's at least as good. And if that world has strictly *greater* value at at some locations, then it's strictly better. This sensitivity to changes in local value seems a minimal requirement for any comparison of worlds which stays true to the spirit of aggregation.

Separability of Value: If W_a and W_b contain the same locations and $W_a \succcurlyeq W_b$ then, adding their corresponding local values, $W_a + W \succcurlyeq W_b + W$ for all $W \in \mathcal{W}$ with the same locations.⁸

This principle, also called *Translation Scale Invariance* in the literature, ensures that betterness is sensitive *only* to differences in local value. If between one pair of worlds there is the same pattern of differences as between another pair of worlds, we must rank both pairs the same way. It does not matter what local values we start with in a world; all that matters is what we add or take away; if a certain combination of additions and removals is an improvement to a world, it would count as an improvement to any world. This is a key distinguishing feature of aggregative theories in the finite context (such as any which endorse totalism or total prioritarianism). And, moving to the infinite context, Separability of Value seems a necessary condition for staying true to the spirit of aggregation.

And the final condition we need, *Finite Sum*, requires only that \succcurlyeq remains consistent with what our judgements would be if the universe would contain only finite total value. I take this as a crucial requirement for extensional adequacy.

Finite Sum: If there is a finite total sum of local values in both W_a and W_b , and the sum in W_a is at least as great as that in W_b , then $W_a \succcurlyeq W_b$.

Each of these principles (and their conjunction) is hard to deny. So Additivity is hard to deny. Given this, it is unsurprising that it is satisfied by every plausible stronger proposals so far proposed (e.g. Vallentyne and Kagan, 1997; Arntzenius, 2014; Jonsson and Voorneveld, 2018; Wilkinson, 2021; Bostrom, 2011, pp. 27-30). Given this broad agreement, I will assume that \succcurlyeq satisfies Additivity. But I remain agnostic about what else it says, as I want my conclusions to hold for all of those stronger views.

⁸Define addition of worlds as follows. For all worlds W_a and W_b with the same locations \mathcal{L} , the world $W_v = W_a + W_b$ is given by $V_v(l) = V_a(l) + V_b(l)$ for all $l \in \mathcal{L}$.

3 Comparing lotteries

We don't just want to compare worlds; we want to compare *lotteries* over those worlds. Formally, a lottery L is represented as a probability measure on (the minimal Boolean algebra containing the elements of) \mathcal{W} — L maps all sets of possible worlds to probabilities in the interval $[0, 1]$, while obeying the standard probability axioms. The set of all such probability measures is denoted by \mathcal{P} . We can also define the *domain* of each lottery L_i by $\mathcal{W}_i = \{W \in \mathcal{W} | L_i(\{W\}) > 0\}$. For two lotteries L_i and L_j , the union of their domains $\mathcal{W}_i \cup \mathcal{W}_j$ can be abbreviated to $\mathcal{W}_{(i,j)}$.

I will make a few other abbreviations to keep the notation in check. $L(\{W\})$ will be abbreviated to $L(W)$ when the input to L is simply $\{W\}$. And $L_i(\succcurlyeq W)$ will be used as an abbreviation for the probability that L_i gives to the set of all worlds in \mathcal{W}_i that are *at least as good as* W (or, equivalently, $\succcurlyeq W = \{W' \in \mathcal{W}_i | W' \succcurlyeq W\}$). And, when a lottery results in just one world W with probability 1, I denote both world and lottery by W .

To compare such lotteries, we need another ‘at least as good as’ relation: a binary relation \succcurlyeq , this time on \mathcal{P} . As before, strict betterness (\succ) and equality (\sim) are represented as the asymmetric and symmetric components, respectively. And again, as basic desiderata, I will assume that \succcurlyeq must be reflexive and transitive. As well, it must be consistent with \succ : if $W_a \succ W_b$ then $W_a \succcurlyeq W_b$ too.

3.1 Local expectations

But what \succcurlyeq relation should we adopt; when should $L_a \succcurlyeq L_b$ hold? One proposal comes from Arntzenius (2014) and Bostrom (2011, pp. 27-30) and is also endorsed by Meacham (2020).

Consider the lotteries: W_{101} , which delivers world W_{101} for sure; and L_1 , which brings even odds of the worlds **2** and **0** (the worlds of value 2 or 0, respectively, at every location).

$$\begin{array}{rcccccccccccc}
 & l_a & l_b & l_c & l_d & l_e & l_f & l_g & l_h & l_i & l_j & l_k & \dots \\
 W_{101} : & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & \dots
 \end{array}$$

$$L_1 \left\{ \begin{array}{l|cccccccccccc}
 L_1(W) & & l_a & l_b & l_c & l_d & l_e & l_f & l_g & l_h & l_i & l_j & l_k & \dots \\
 1/2 & \mathbf{2} : & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & \dots \\
 1/2 & \mathbf{0} : & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \dots
 \end{array} \right.$$

If Additivity holds, then $\mathbf{2} \succ W_{101} \succ \mathbf{0}$. But it says nothing about W_{101} and L_1 as lotteries. Nor does Additivity alone allow us to construct expected values: the total differences in value between these worlds are infinite, so there’s no clear way to assign real values to outcomes over which we can calculate expectations.

Without the opportunity to apply expected value theory, Arntzenius, Bostrom, and Meacham recommend we do the following. Consider the prospects of each individual location. In L_1 , each location has probability $\frac{1}{2}$ of 2, and probability $\frac{1}{2}$ of 0. So we can take the expected *local* value for each location l_i .

$$\begin{array}{cccccccccccccc}
 & l_a & l_b & l_c & l_d & l_e & l_f & l_g & l_h & l_i & l_j & l_k & \dots \\
 \mathbb{E}_{L_1}(V(l_i)) : & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & \dots
 \end{array}$$

By doing so, we can define a new object with the same structure as a world—an ‘expected world’ $\mathbb{E}W_L$, with the same locations and with value function given by the *expected* local values under the lottery. Equivalently, $\mathbb{E}W_L = \langle \mathcal{L}, \mathbb{E}_L(V) \rangle$.

This expected world can easily be compared to W_{101} , via Additivity (or even by Pareto alone). And it’s strictly better. So we might conclude that $L_1 \succ W_{101}$.

This approach can be stated as *Local Expecations* which, to an approximation, is endorsed by each of Arntzenius⁹, Bostrom¹⁰, and Meacham. (See notes for differences.)

⁹Arntzenius (2014, pp. 55-6) proposes the following.

Weak Location Criterion: For lotteries A and B , $A \succ B$ iff $\sum_{l \in \mathcal{L}} (\mathbb{E}_A(V(l)) - \mathbb{E}_B(V(l)))$ “...is absolutely convergent and > 0 , where we are summing over all (epistemically possible) [locations]...” $l \in \mathcal{L}$.

This is *almost* equivalent to the conjunction of Local Expectations and Additivity. Except: 1) it defines only a strict betterness relation \succ , so does not imply that $L_1 \sim L_2$ for any lotteries; 2) it remains silent if the sum diverges unconditionally to $+\infty$, even in cases of certainty in which Additivity gives a verdict; and 3) it allows worlds in both lotteries to contain different locations (e.g., the same people) summing value instead over all epistemically possible locations, which requires that we assign some value to the non-existent lives. I don’t want to make a stand on such cases here, and my modifications to (1) and (2) will likely be uncontroversial.

¹⁰Bostrom’s (2011: pp. 27-30) proposal is more complicated. He describes (but doesn’t explicitly endorse) an approach by which we represent each world’s total value with a *hyperreal number*: a vector of (countably) infinite length that consists of the cumulative sums of local values, summed in some common order. In the example above, if we chose to sum in the order l_a, l_b, l_c , etc, then W_{101} ’s total would be represented by the hyperreal $(1, 1, 2, 2, 3, 3, \dots)$ and $\mathbf{2}$ ’s by $(2, 4, 6, 8, 10, \dots)$. For one world to be better than another, the entries in its hyperreal total must be greater than the those in the other at ‘sufficiently many’ positions. For instance, $\mathbf{2}$ has larger entries than W_{101} at *all* positions, so it would be better.

The standard of ‘sufficiently many’ here can vary, but one very minimal condition is that, if the entries of one hyperreal are greater than another in all but finitely many positions, then the world associated with the former is better. In effect, this condition is equivalent to Additivity. (There are stronger conditions we might apply too, but those can be chosen to make the hyperreal approach equivalent to any plausible aggregation rule we want—see Pivato (2014).)

Local Expectations: For any lotteries $L_a, L_b \in P$ such that all $W \in \mathcal{W}_{(a,b)}$ contain the same locations, $L_a \succcurlyeq L_b$ if $\mathbb{E}W_{L_a} \succcurlyeq \mathbb{E}W_{L_b}$.

If both Local Expectations and Additivity hold, then we can make the comparison between the above lotteries W_{101} and L_1 , as demonstrated. And, in general, we then have the result that $L_a \succcurlyeq L_b$ if the following sum is (unconditionally) greater than or equal to 0 (or diverges unconditionally to $+\infty$).¹¹

$$\sum_{l \in \mathcal{L}} \left(\mathbb{E}_{L_a}(V(l)) - \mathbb{E}_{L_b}(V(l)) \right)$$

This seems promising: it appears we can sidestep the whole problem of taking expectations over infinite totals. We just need to lower our expectations, down to the local level.

4 The problem

But Local Expectations has a problem. Consider the case of *Egregious Energy*.

Example: Egregious Energy

A new energy source has been discovered, and you must decide whether humanity takes advantage of it.

If we do, we will produce enormous amounts of energy and many lives will be improved. But there are a few downsides. One is that the fuel needed is limited—we will only reap the benefits for only a short time. Another is that it may take some time to get working but, the longer it takes, the longer that cornucopia of energy will last. But the greatest problem is pollution—this energy source produces a novel form of pollution which will badly harm human health. That pollution and its effects will decrease over time, but we will never fully eradicate it; it will continue to harm human well-being for the entire future of humanity (which I’ll assume will be infinitely long). In short, using this energy source will produce some finite benefit, but cause infinite total harm. (The relevant

We can also sum hyperreals and multiply them by real numbers (just as is done with vectors), so they can give us expected values in roughly the old-fashioned way. L_1 would have expected value $(1, 2, 3, 4, 5, 6, \dots)$, which is precisely the hyperreal total of $\mathbb{E}W_{L_1}$. Since this is larger than W_1 ’s total, we could say that $L_1 \succ W_{101}$. In effect, this is equivalent to applying Local Expectations.

¹¹Combined with Additivity or any other ‘at least as good’ relation over worlds, Local Expectations gives us a strengthened form of what is often called *ex ante* Pareto.

probabilities and values are given below.)

For simplicity, exactly the same persons (and person-time-slices, and other possible types of locations) will exist at exactly the same physical positions whether or not we adopt this new energy source.

Since the same persons exist at the same physical positions either way, we can treat the resulting worlds as having the same locations. Then we can represent your options as the world $\mathbf{0}$ and the lottery L . If you forego the energy source, you produce $\mathbf{0}$, simply representing the baseline of what would have happened otherwise.¹² And if you choose to exploit it, you produce L , a lottery over infinitely many worlds W_1, W_2, \dots, W_j , and so on (each with probability $\frac{1}{2^j}$). You're uncertain of how long it takes the energy source to start working, during which time everyone obtains the same value 0 as in $\mathbf{0}$. Once it's working, some number of people obtain greater value, represented by 2. And then, once the fuel runs out, every person obtains less value than the baseline, represented by a negative number. Note that, in every world in L , the total value diverges unconditionally to $-\infty$. So they're all *far* worse than $\mathbf{0}$.¹³ Note that the sequence of locations (l_1, l_2, l_3, \dots) is chronological, but nothing hangs on this.

		l_1	l_2	l_3	l_4	l_5	l_6	l_7	l_8	\dots		
	$\mathbf{0} :$	0	0	0	0	0	0	0	0	\dots		
L	{	$L(W)$	l_1	l_2	l_3	l_4	l_5	l_6	l_7	l_8	\dots	
		1/2	$W_1 :$	2	-1/2	-1/2	-1/4	-1/4	-1/4	-1/4	-1/8	\dots
		1/4	$W_2 :$	0	2	2	-1/4	-1/4	-1/4	-1/4	-1/8	\dots
		1/8	$W_3 :$	0	0	0	2	2	2	2	-1/8	\dots
		1/16	$W_4 :$	0	0	0	0	0	0	0	2	\dots
		\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots		

¹²Note that local values of $\mathbf{0}$ here do not imply that the lives in question are right on the boundary of not worth living. They might be extremely valuable lives. But these local values are here represented cardinally—the numbers only capture the relative size of the differences between them. So these same representations of $\mathbf{0}$ and L could describe worlds in which everyone has a blissful life and suffers only a slight reduction in quality of life due to the pollution. Or they could describe worlds in which everyone suffers terribly and that pollution makes life even worse.

¹³In each world W_i , the local value at l_j is given by:

$$V_i(l_j) = \begin{cases} 0 & \text{for } i < 2^j \\ 2 & \text{for } 2^i \leq j < 2^{i+1} \\ -\frac{1}{2^k} & \text{for } 2^k \leq j < 2^{k+1}, \forall k > i \end{cases}$$

We can compare L to $\mathbf{0}$ using Local Expectations. As above, take L 's expected value for each location to obtain $\mathbb{E}W_L$.

$$\begin{array}{cccccccccc} & l_1 & l_2 & l_3 & l_4 & l_5 & l_6 & l_7 & l_8 & \cdots \\ \mathbb{E}W_L : & 1 & 1/4 & 1/4 & 1/16 & 1/16 & 1/16 & 1/16 & 1/64 & \cdots \end{array}$$

The expected local values are all greater than 0, so Additivity (or just Pareto alone) implies that $\mathbb{E}W_L \succ \mathbf{0}$. Together with Local Expectations, this implies that $L \succ \mathbf{0}$. It is allegedly better to adopt the energy source.

But I would suggest that this verdict is implausible. After all, L guarantees us a worse outcome—*every one of its possible outcomes* is worse than $\mathbf{0}$ (by Additivity). Yet Local Expectations still says that it is the better choice. This clashes with my own intuitions, and is a violation of *Guaranteed Betterness*.

Guaranteed Betterness: If *every* world in \mathcal{W}_a (the domain of L_a) is better than *every* world in \mathcal{W}_b , then $L_a \succ L_b$.

This conflict can be stated more formally as Theorem 1.¹⁴

Theorem 1: For any reflexive, transitive relation \succ on \mathcal{P} , if \succ satisfies Local Expectations then it cannot satisfy both Guaranteed Betterness and Additivity.

Is this so bad? I certainly think it is. Although Local Expectations seems plausible, Guaranteed Betterness seems to me *undeniable*. Why? The first reason is raw intuition—in this case, and in other cases where the two conflict, my own intuition tells me that it is clearly worse to make a lottery's outcome worse with certainty. But some might have conflicting intuitions. For some, it may be more intuitively plausible to reject Guaranteed Betterness than it is to recommend making every single individual's prospects worse. Intuition may more strongly favour particular judgements in (at least some) cases like this than it does the broad principle of Guaranteed Betterness.

For readers not immediately convinced of Guaranteed Betterness, I'll offer two arguments beyond mere intuition. This first is this. A crucial desideratum of any decision theory—indeed, I think,

¹⁴To prove this would be straightforward. In comparing $\mathbf{0}$ to L , Guaranteed Betterness and Additivity together imply that W_0 is strictly better, while Local Expectations and Additivity together imply that L is.

the *most* crucial desideratum—is that its recommendations help us obtain what we ultimately care about. As Schoenfield (2014, p. 268) points out, any decision theory “...that makes demands that don’t make sense given our concern with *value* can’t do what [a decision theory] is meant to do...” and so should be rejected. And when applying an aggregative theory to make decisions for a large population, we ultimately care about bringing about outcomes that have the most value *overall*, not just the most value (or best prospects) for any particular individual/s. If a decision theory recommends lotteries that are sure not to help us bringing about better outcomes (or, even worse, make the resulting outcome worse) then it fails on this desideratum. And Local Expectations clearly fails here. It judges that L is better than $\mathbf{0}$ even though it is guaranteed to produce a worse outcome. It thus clashes with our basic goal of producing better outcomes so, I claim, we must reject it.

Another argument for Guaranteed Betterness is as follows. Moral betterness is more fundamental than instrumental moral betterness—the ranking of lotteries is grounded in the ranking of worlds by their overall value, not the other way around. If a world is better than another it is due simply to the good-making features of each; not due to the fact that the lottery with probability 1 of one of those outcomes is better than the other corresponding lottery. But, if we adopt Local Expectations and reject Guaranteed Betterness, we are committed to instrumental moral betterness being more fundamental. To see why, note that in Egregious Energy we could replace each world in L for another world that is equally good according to any ranking that satisfies Additivity: replace W_1 , which had local values $(2, -\frac{1}{2}, -\frac{1}{2}, -\frac{1}{4}, -\frac{1}{4}, -\frac{1}{4}, -\frac{1}{4}, -\frac{1}{8}, \dots)$ for locations listed in the same order as above, with a world with local values $(0, 0, 0, 0, 0, 0, 0, -\frac{1}{8}, \dots)$; so too, replace each other world W_i with one where all positive value is redistributed to later locations so that all local values are either 0 or negative. The expected local value for each location is then 0 or negative, and so Local Expectations will say that this new lottery is worse than $\mathbf{0}$. But this cannot happen if rankings of lotteries are grounded in rankings of worlds—any plausible ranking of worlds says that W_1 and its replacement are equally good and each compare to every other world in the same way, and likewise for each other world and its replacement. So, by Local Expectations, the moral ranking of these worlds does not fully determine the ranking of lotteries, and that is implausible.

In my view, even stronger requirements are placed on our decision theory here by the desideratum that it help us to obtain what we ultimately care about. Suppose that one lottery L_a has at least as great a probability of turning out at least as good as some world W as some other lottery L_b . And suppose that this holds for *all* worlds W —for *any* world, L_a has as high or higher probability

of resulting in W or something even better. If our goal is to promote value, surely L_a must be at least as good a lottery as L_b . So says *Stochastic Dominance*.

Stochastic Dominance: Let $L_a, L_b \in \mathcal{P}$ be such that $\mathcal{W}_{(a,b)}$ is totally ordered by \succcurlyeq . If $L_a(\succcurlyeq W) \geq L_b(\succcurlyeq W)$ for all $W \in \mathcal{W}_{(a,b)}$, then $L_a \succcurlyeq L_b$.

If, as well, $L_a(\succcurlyeq W) > L_b(\succcurlyeq W)$ for some $W \in \mathcal{W}$, then $L_a \succ L_b$.

This is a strengthening of Guaranteed Betterness, but not a radical one. When dealing with finite payoffs, both principles are consistent with, but weaker than, standard expected value theory—for instance, they do not rule out risk aversion (see Buchak, 2013). Accept expected value theory and we must accept Stochastic Dominance. But, while expected value theory is somewhat controversial in the existing literature, Stochastic Dominance and its kin are not. Only in rare circumstances have philosophers proposed normative decision theories that violate this form of Stochastic Dominance.¹⁵

But Local Expectations violates both Stochastic Dominance and the far weaker Guaranteed Betterness. So, if we wish to retain either principle, we must reject Local Expectations. But if we do so, how then can we compare lotteries over infinite worlds?

5 My proposal

I will propose two methods of comparing such lotteries: one weaker and able to handle easy cases; and one stronger and able to handle more difficult cases.

5.1 Expectations of Differences 1

I'll start with a class of easy cases: comparing lotteries in which each pair of worlds differ by at most a finite sum of local differences. L_2 and $\mathbf{1}$ are two such lotteries.

$$L_3 \left\{ \begin{array}{l|lllllll} L_3(W) & & l_a & l_b & l_c & l_d & l_e & l_f & \cdots \\ \hline 1/2 & W_5 : & 2 & 2 & 2 & 1 & 1 & 1 & \cdots \\ 1/2 & W_0 : & 0 & 0 & 1 & 1 & 1 & 1 & \cdots \end{array} \right.$$

¹⁵Examples include Seidenfeld et al. (2009); Smith (2014) and Lauwers and Vallyntyne (2016). Schoenfield (2014) rejects a similar, but distinct, form of Stochastic Dominance which gives verdicts even when outcomes in the domains aren't totally ordered by the betterness relation. She raises no objection to weaker formulations like mine.

$$\begin{array}{cccccccc}
& l_a & l_b & l_c & l_d & l_e & l_f & \cdots \\
\mathbf{1} : & 1 & 1 & 1 & 1 & 1 & 1 & \cdots
\end{array}$$

Here's how we might proceed. First, with just Additivity, we can rank the worlds $W_5 \succ \mathbf{1} \succ W_0$.

Next, we might pick a world as a baseline—say, W_0 —and represent each world by the sum of differences between it and W_0 . For W_0 , that's 0. For W_5 , it's 5. And for $\mathbf{1}$, it's 2. So we have a nice cardinal 'total' for each world. And that gives us all we need to calculate expected 'total' values.

For example:

$$L_2(W_0) \cdot 0 + L_2(W_5) \cdot 5 = \frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 5 = 4\frac{1}{2}$$

Calculated this way, the expected 'total' of L_3 is less than that of $\mathbf{1}$ (which is just 2), so we might claim that L_2 is better. Putting this approach more precisely, we have *Expectations of Differences 1*.

Expectations of Differences 1 (ED1): For any $L_a, L_b \in \mathcal{P}$ such that all worlds in $\mathcal{W}_{(a,b)}$ contain the same locations \mathcal{L} , if there exists a world $W_* \in \mathcal{W}_{(a,b)}$ such that the sum

$$\sum_{W_i \in \mathcal{W}_a} L_a(W_i) \left(\sum_{l \in \mathcal{L}} (V_i(l) - V_*(l)) \right)$$

is at least as great as the corresponding sum for L_b , then $L_a \succcurlyeq L_b$.

ED1 differs slightly from how we usually calculate finite expectations. In the finite setting, we usually take the expectation of the *total* value in each world and see how they differ—effectively, we consider the *differences of expectations*. But my approach is to consider the *expectations of differences*. We represent the value of each world by the sum of its differences from some 'baseline' world W_* . And then, for each lottery, we take the expectation of that sum. If we wanted, we could do this in finite cases and reach the same verdicts. But in infinite cases, these approaches come apart, and the latter clearly does better.

And ED1 doesn't just work when the sums of differences between worlds are finite. It can also work when some of those sums diverge unconditionally to (positive or negative) infinity. As long as the worlds with positively infinite sums all appear in one lottery and the worlds with negatively infinite sums appear in the other, we can still say which lottery is better. For instance, recall the

lotteries from Egregious Energy.

$$\begin{array}{cccccccccc}
 & l_1 & l_2 & l_3 & l_4 & l_5 & l_6 & l_7 & l_8 & \cdots \\
 \mathbf{0} : & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots \\
 \\
 L \left\{ \begin{array}{l} L(W) \\ 1/2 \\ 1/4 \\ 1/8 \\ 1/16 \\ \vdots \end{array} \right. & \left| \begin{array}{l} \\ W_1 : \\ W_2 : \\ W_3 : \\ W_4 : \\ \vdots \end{array} \right. & \begin{array}{cccccccccc}
 l_1 & l_2 & l_3 & l_4 & l_5 & l_6 & l_7 & l_8 & \cdots \\
 2 & -1/2 & -1/2 & -1/4 & -1/4 & -1/4 & -1/4 & -1/8 & \cdots \\
 0 & 2 & 2 & -1/4 & -1/4 & -1/4 & -1/4 & -1/8 & \cdots \\
 0 & 0 & 0 & 2 & 2 & 2 & 2 & -1/8 & \cdots \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2 & \cdots \\
 \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots
 \end{array}
 \end{array}$$

ED1 implies that $\mathbf{0}$ is better than L , unlike the Arntzenius-Bostrom-Meacham approach. First, let the ‘baseline world’ be $\mathbf{0}$. Then the sum of differences between $\mathbf{0}$ and $\mathbf{0}$ is simply 0. Meanwhile, each world in L has a sum of differences from $\mathbf{0}$ that diverges unconditionally to $-\infty$.¹⁶ But since all of these worlds with infinite sums appear in one lottery rather than the other, ED1 can say which lottery is better. The expected sum of differences for L will diverge unconditionally to $-\infty$, while the expected sum for $\mathbf{0}$ will be 0 (a lot greater than $-\infty$!). So, by ED1, $\mathbf{0} \succ L$, *contra* the Arntzenius-Bostrom-Meacham approach.

ED1 has other arguments in its favour too. Intuitively, it seems a plausible and natural way to judge lotteries. It’s a rough analogue of Additivity for this new setting of comparing lotteries. In fact, in cases of certainty, it *implies* Additivity. And, like Additivity, it is implied by some highly plausible conditions. (All proofs appear in the appendix.)

Theorem 2: For any reflexive, transitive relation \succsim on \mathcal{P} , if \succsim satisfies Stochastic Dominance, Finite Expectations, and Separability of Value (for Lotteries) then it satisfies

¹⁶The calculation is:

$$\begin{aligned}
 & \sum_{W_i \in \mathcal{W}} L(W_i) \left(\sum_{l \in \mathcal{L}} (V_i(l) - V_0(l)) \right) \\
 &= \sum_{W_i \in \mathcal{W}} \frac{1}{2^i} \left(\sum_{l \in \mathcal{L}} (V_i(l) - 0) \right) \\
 &= \sum_{W_i \in \mathcal{W}} \frac{1}{2^i} \left(0 + 0 + \dots + 2 \cdot 2^{i-1} + \sum_{k=1}^{\infty} -1 \right) \\
 &\rightarrow -\infty
 \end{aligned}$$

ED1.

Stochastic Dominance will be familiar from the previous section. Recall that it is consistent not just with standard expected value theory but also with risk aversion. To impose risk neutrality, I use *Finite Expectations*. This is the analogue of Finite Sum (from Section 2) for comparing lotteries. It requires only that \succsim remains consistent with what our judgements would be if our lotteries contained only finite expected total value. If \succsim doesn't satisfy this, then it isn't an adequate extension of finite expected value theory.

Finite Expectations: For any $L_a, L_b \in \mathcal{P}$ such that each lottery has finite expected total sum of local values k_a or k_b , respectively, $L_a \succsim L_b$ if and only if $k_a \geq k_b$.

Then there's the highly plausible *Separability of Value (for Lotteries)*. Much like Separability of Value from earlier, it says that we can take any two lotteries and add whatever string of local values we want to every world in both, and the resulting lotteries will be ranked the same way. This implies that comparisons of lotteries are sensitive *only* to differences in the probabilities and local values of the worlds in their domains, and that those local values are additively separable. This is the direct analogue of Separability of Value for the lottery context, such that I will abbreviate it to 'Separability of Value' in what follows.

Separability of Value (for Lotteries): For any $L_a, L_b \in \mathcal{P}$ such that all worlds in $\mathcal{W}_{(a,b)}$ contain the same locations, and for any $W \in \mathcal{W}_{(a,b)}$, let L'_a denote the lottery with $L'_a(W) = L_a(W - W')$ for all $W \in \mathcal{W}_a + W'$, and similarly for L'_b . If $L_a \succsim L_b$ then $L'_a \succsim L'_b$.¹⁷

I find Stochastic Dominance, Finite Expectations, and Separability of Value, as well as their conjunction, hard to deny. So ED1 is on firm ground as a minimal principle for comparing lotteries, just as Additivity is for comparing worlds.

5.2 Expectations of Differences 2

But ED1 doesn't get us far. Consider W_{101} and L_3 .

¹⁷This also implies the converse since, for all W' , the condition also applies to $-W'$.

$$\begin{array}{cccccc}
& l_a & l_b & l_c & l_d & l_e & \cdots \\
W_{101} : & 1 & 0 & 1 & 0 & 1 & \cdots
\end{array}$$

$$L_3 \left\{ \begin{array}{l|cccccc}
L_3(W) & & l_a & l_b & l_c & l_d & l_e & \cdots \\
\frac{1}{2} & W_3 : & 3 & 3 & 3 & 3 & 3 & \cdots \\
\frac{1}{2} & W_0 : & 0 & 0 & 0 & 0 & 0 & \cdots
\end{array} \right.$$

Take any pair of those worlds and sum their local differences; the result is infinite. If we apply ED1 here, no matter which baseline-world we pick, either both expected sums are infinite or at least one is undefined. Using ED1 alone, we cannot compare these lotteries. Yet, intuitively, we can. So a stronger rule would be helpful.

Before I present that rule, consider how little we need to compare two lotteries over finite payoffs. We can make judgements without knowing the probabilities in each lottery—we just need to know the *differences* in probability of each outcome between one lottery and the other. And we can make judgements without knowing the values of the outcomes—we just need to know the (scale of the) differences between them. Even then, we need not know the *precise* differences—upper and lower bounds can be enough.

To illustrate, here are two lotteries over finite payoffs that we can compare with incomplete information.

L'_{101} : value 1 with probability 1.

L'_3 : value $v + 1$ with probability $\frac{1}{2}$; value 0 otherwise.

Here, v is some real number such that $v > 2$. We cannot assign it a precise value, perhaps because it is vague or indeterminate. But we can still judge which lottery is better. Take their expected values: L'_2 has expected value 1; and L'_3 has expected value $\frac{v}{2}$, which is greater than 1. So L'_3 is better. Incomplete information is unproblematic here.

We have similar information when comparing the above lotteries W_{101} and L_3 over infinite worlds—they are structurally equivalent to L'_2 and L'_3 . So we can take a similar approach.

First, take the difference between each world and the next best world in the domain of the two lotteries. If our worlds had finite total values, we would represent those differences with finite values.

But, with infinite worlds, we must represent differences as worlds themselves—worlds given by the differences in local values. For each world W_i which has some distinct next best world W_j in the domain, the difference between them is given by $D_i = \langle \mathcal{L}, V_i - V_j \rangle$, with local values $V_i(l) - V_j(l)$ at every location. (Note that, if there is no world in $\mathcal{W}_{(1,2)}$ that is strictly worse than W_i , then D_i is undefined.)

$$\begin{array}{cccccc} & l_a & l_b & l_c & l_d & l_e & \cdots \\ D_3 : & 2 & 3 & 2 & 3 & 2 & \cdots \\ D_2 : & 1 & 0 & 1 & 0 & 1 & \cdots \end{array}$$

Second, compare those differences to one another. Again, this would be straightforward with finite differences, but less so here. But, fortunately, we *can* compare them. Those differences are themselves worlds, and can be compared using the same ‘at least as good as’ relation as we would use for any worlds. Since that relation obeys Additivity (or just Pareto), $D_3 \succ D_1$.

Further, we can say *how much better* they are than one another by taking scalar multiples of them.¹⁸ A scalar multiple of a world W can be defined as $k \cdot W = \langle \mathcal{L}, k \times V_i \rangle$ for any real k —a world with the same locations but local values k times as great. For instance,

$$\begin{array}{cccccc} & l_a & l_b & l_c & l_d & l_e & \cdots \\ 2 \cdot D_1 : & 2 & 0 & 2 & 0 & 2 & \cdots \end{array}$$

Note that D_3 is still greater than $2 \cdot D_2$. So we can say that D_3 is more than twice as great a difference as D_2 . How much greater, exactly? It doesn’t matter, to compare these lotteries. The fact that $D_3 \succ 2 \cdot D_2$ gives us all the information we need. With that fact in hand, the information we have here is analogous to what we had for the analogous finite lotteries L'_2 and L'_3 above. D_3 is analogous to v and D_2 to the value 1. So we might compare these lotteries in an analogous way, and say that $L_3 \succ W_{101}$.

Returning to the finite context, more generally, we do not even need to know the exact probabilities involved in lotteries like L'_2 and L'_3 . It suffices to know how much greater they are for one lottery than another. This is because the statement that some lottery L_a has greater expected value

¹⁸Since the local values in the original worlds W_i and W_j are represented on a common interval scale, the local values in D_i can be represented on a *ratio scale*—they have an absolute zero, 0. So it makes sense to compare their absolute size via scalar multiplication, and likewise for such difference-worlds at large.

than some L_b can be rewritten as:

$$\mathbb{E}(L_a) - \mathbb{E}(L_b) = \sum_{w \in \mathcal{W}_{(a,b)}} \left((L_a(w) - L_b(w)) \times V(w) \right) \geq 0$$

Further, it can be rewritten in terms of differences in probabilities *and* differences in value, as below. Here, we use the probability of each lottery turning out at least as good as each outcome w_i , rather than the probability of w_i alone. (Note that the outcomes w_1, w_2, w_3, \dots are ordered from worst to best.)

$$\mathbb{E}(L_a) - \mathbb{E}(L_b) = \sum_{w_i \in \mathcal{W}} \left((L_a(\succsim w_i) - L_b(\succsim w_i)) (V(w_i) - V(w_{i-1})) \right) \geq 0$$

In the infinite context, I propose that we attempt to satisfy something akin to this equation—evaluating the expectations of differences—rather than taking the standard approach of taking expectations and comparing them. I propose that we apply *Expectations of Differences 2*, which goes like this.

Expectations of Differences 2 (ED2): Take any $L_a, L_b \in \mathcal{P}$ such that all $W \in \mathcal{W}_{(a,b)}$ have the same locations. $L_a \succcurlyeq L_b$ if, for some $W_* \in \mathcal{W}_{(a,b)}$ and for all $W_i \in \mathcal{W}_{(a,b)}$, there exists $k_i \in \mathbb{R}$ such that

$$\sum_{W_i \in \mathcal{W}_{(a,b)}} \left(L_a(\succsim W_i) - L_b(\succsim W_i) \right) k_i \geq 0$$

$$\text{and either } k_i \cdot D_* \preccurlyeq D_i \quad \text{if } L_a(\succsim W_i) - L_b(\succsim W_i) > 0,$$

$$\text{or } k_i \cdot D_* \succcurlyeq D_i \quad \text{if } L_a(\succsim W_i) - L_b(\succsim W_i) < 0.$$

The first equation here is an analogue of the equation above, but k_i stands in as a measure of just how great the difference is between W_i and the next best world in the domain W_j . And the later conditions ensure that k_i is such a measure. Compared to some baseline difference-world D_* , k_i is what we must multiply D_i by for it to be greater (or smaller, if getting W_i or better is more likely under L_b). But k_i need only be a rough measure: if L_a has the higher probability of producing W_i or better, then k_i is just a lower bound on the size of D_i ; if L_b has the higher probability of W_i or better, then k_i serves as an upper bound on the size of D_i . We don't need k_i to give exact relative sizes.

If we apply ED2 to the comparison of W_{101} and L_3 from earlier, the process matches the informal reasoning given above. With only three worlds in the domain $\mathcal{W}_{(2,4)}$, we have just two difference-

worlds D_3 and D_2 (as detailed above). We can treat D_2 as the baseline D_* . Then k_2 can be 1 (since $1 \cdot D_2 \sim D_2$), k_3 can be 2 (since $2 \cdot D_3 \succ D_2$), and k_0 can be undefined as D_0 is. These k_i s satisfy the latter conditions of ED2. And when we move to the former condition, the sum on the left of the equation becomes $\frac{1}{2} \times 2 + 0 \times 1 + 0 \times k_0 = 1 > 0$. So ED2 is satisfied; L_3 is better than W_{101} , as intuition suggests.¹⁹

But should we accept ED2? I think so, and not just based on the intuitive correctness of its verdicts or its analogy to finite decision-making. For one, it implies all of the (very plausible) judgements of ED1, as long as Additivity holds.

Theorem 3: For any reflexive, transitive relation \succsim on \mathcal{P} , if \succsim satisfies Expectations of Differences 2 and Additivity, then it satisfies Expectations of Differences 1.

Given the compatibility of ED2 with ED1, we already know that this version of Expectations of Differences will deal with cases like Egregious Energy better than rival views. But it's also stronger than ED1, as demonstrated by the example above.

Also in its favour, the rule is implied by the conjunction of several highly plausible principles.

Theorem 4: For any reflexive, transitive relation \succsim on \mathcal{P} , if \succsim satisfies Stochastic Dominance, Separability of Value (for Lotteries), Independence, and Extrapolated Expectations, then it satisfies Expectations of Differences 2.

Stochastic Dominance and Separability of Value will be familiar from earlier. Then we have two newcomers.

First, *Independence* is the same principle that is often used to axiomatise expected utility theory. Suppose we evaluate L_1 is at least as good as L_2 . Independence says that we could mix each lottery with some third lottery L_3 , whatever it might be, and the resulting mixed lotteries would be ranked the same way. More formally, it can be stated as follows.

Independence: For any $L_a, L_b, L_c \in \mathcal{P}$ define lotteries $L_{a \vee c}, L_{b \vee c}$ as

$$L_{a \vee c}(W) = p \times L_a(W) + (1 - p) \times L_c(W) \quad \text{for all } W \in \mathcal{W}$$

¹⁹Note that it is only strictly better because there exist no suitable k_3, k_2 , and k_0 that make this sum less than or equal to 0.

$$\text{and } L_{b\vee c}(W) = p \times L_b(W) + (1 - p) \times L_c(W) \quad \text{for all } W \in \mathcal{W}$$

For any $L_a, L_b, L_c \in \mathcal{P}$ and for any $p \in [0, 1]$, $L_{a\vee c} \succcurlyeq L_{b\vee c}$ if and only if $L_a \succcurlyeq L_b$.

We then have *Extrapolated Expectations*, which should also be uncontroversial for those who accept expected value theory in the finite setting. It merely says that, for any world $W' \succ \mathbf{0}$, a lottery with probability p of W' (and $p - 1$ of $\mathbf{0}$) is precisely as good as $p \cdot W'$.

Extrapolated Expectations: For any world $W' \succ \mathbf{0}$ and any $p \in (0, 1]$, let L be defined as below. Then $L \sim p \cdot W'$.

$$L(W) = \begin{cases} p & \text{for } W = W' \\ 1 - p & \text{for } \mathbf{0} \\ 0 & \text{for } W \notin \{W', \mathbf{0}\} \end{cases}$$

If we faced a similar lottery over finite payoffs, w and 0, we'd assign it expected value $p \cdot w$ without hesitation. Extrapolated Expectations requires that we treat infinite worlds in the same way. And it's on firm footing of its own—if our \succcurlyeq relation for comparing worlds obeys Additivity, then Extrapolated Expectations follows straightforwardly from Finite Expectations and Stochastic Dominance. It is also worth noting that it follows from Local Expectations—if one is tempted to adopt Local Expectations instead of ED2, Extrapolated Expectations must hold either way.

And if we accept Extrapolated Expectations, Independence, Stochastic Dominance, and Separability of Value, then we must also accept ED2, by Theorem 4.

5.3 Weaknesses of Expectations of Differences

Even ED2 does not provide judgements in *every* case. For one, the rule's antecedent requires that *every* world in the domain of both lotteries contains the same locations. For another, to assign a suitable k_i to each difference-world, many of these differences need to be comparable to certain scalar multiples of one another. And these conditions sometimes are not met.

Here is one pair of lotteries where the latter condition isn't met. They are similar to those I used earlier to demonstrate Local Expectations in action, but with *slightly* different probabilities. (Here, ϵ is some small positive number.)

$$\begin{array}{cccccccccccc}
& l_a & l_b & l_c & l_d & l_e & l_f & l_g & l_h & l_i & l_j & l_k & \dots \\
W_2 : & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & \dots
\end{array}$$

$$L_4 \left\{ \begin{array}{l|cccccccccccc}
L_4(W) & & l_a & l_b & l_c & l_d & l_e & l_f & l_g & l_h & l_i & l_j & l_k & \dots \\
1/2 - \epsilon & W_2 : & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & \dots \\
1/2 + \epsilon & W_0 : & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \dots
\end{array} \right.$$

What does ED2 say? Well, here are our differences-below, and a relevant scalar multiple.

$$\begin{array}{cccccc}
& l_a & l_b & l_c & l_d & l_e & \dots \\
D_2 : & 1 & 2 & 1 & 2 & 1 & \dots \\
D_{101} = W_{101} : & 1 & 0 & 1 & 0 & 1 & \dots \\
(1 + \epsilon') \cdot D_{101} : & 1 + \epsilon' & 0 & 1 + \epsilon' & 0 & 1 + \epsilon' & \dots
\end{array}$$

By Additivity alone, we cannot compare D_2 to the scalar multiple $(1 + \epsilon') \cdot D_{101}$, since the former does better at locations l_b, l_d, l_f, \dots and the latter does better at the rest. The sum of their local differences is undefined. And given the probabilities here, there is no way to generate k_i s that satisfy both of the equations needed for ED2.

Thus, ED2 and Additivity together say nothing about how we should compare W_1 to L . And this is *even though* every pair of worlds at play is comparable and, on top of that, so are the differences between each of them. Even worse, we only made the smallest of changes to a lottery that would otherwise be comparable to W_2 —all we did was give L_4 a mere $+\epsilon$ of probability mass for one outcome, and now we cannot say a thing!

My first response to this silence is simply: *tu quoque*. Suppose we adopt the rival Arntzenius-Bostrom-Meacham view, and so swap ED2 for Local Expectations. Then we still cannot compare these two lotteries. So I am doing no worse here. But that is little comfort.

My second response is that this is not a shortcoming of Expectations of Differences, but of Additivity. Additivity is a very weak constraint on betterness. I have used it so far *because* it is so weak—this weakness makes it uncontroversial and, indeed, all of the plausible stronger principles in the literature are consistent with it. We can treat the conjunction of ED2 and Additivity similarly: it is just a weak and (hopefully) uncontroversial condition for comparing lotteries. So it is fitting

that it makes no judgement in this case. The correct judgement is not obvious, so it should remain silent.

My third, closely related, response is that, actually, ED2 allows us to do better. It can be combined with almost any betterness relation we choose—the definition above makes no reference to Additivity, but instead to some unspecified \succsim relation. We might adopt the \succsim relation from Vallentyne and Kagan (1997, p. 19), Jonsson and Voorneveld (2018), or Wilkinson (2021), each of which is much stronger than Additivity. If in the above example the locations have the right structure, each of these proposals can easily say that L_4 is strictly better than W_{101} .²⁰

And even without a betterness relation stronger than Additivity, ED2 does just fine in the examples in the previous sections, including Egregious Energy. So this approach is not so weak after all.

6 Conclusion

For those who endorse moral theories that rely on an aggregative theory of value, discovering that the universe is infinite may be cause for dismay. Such theories seem to fall silent in all cases where it must rely on facts of betterness. In effect, they seem to say that it is impossible to make the world better.

There is some cause for relief: we have proposals for betterness relations that resemble finite aggregationism but which fare better in infinite worlds (e.g. Vallentyne and Kagan, 1997; Lauwers and Vallentyne, 2004; Arntzenius, 2014; Jonsson and Voorneveld, 2018; Wilkinson, 2021; Bostrom, 2011, pp. 27-30). But these are of little use to limited epistemic agents like us—we are uncertain of the effects of our actions. After all, most of these betterness relations offer no clear way to compare *lotteries* over infinite outcomes. And those that do so have implausible implications, as we saw in Section 4.

This would seem to be a dismal situation for aggregative theories. But I have shown that the situation is rather less dismal than this. No matter which of the above proposals for comparing worlds we endorse, we now have a plausible method for extending that proposal to compare lotteries as well. Just how wide a range of lotteries we can compare will depend on the strength of our method

²⁰In particular, if we combine ED2 with the proposal in Wilkinson (*ibid.*, p. 1946), it will deliver verdicts in almost every physically realistic decision scenario ever faced by agents like us (see Author n.d.).

for comparing worlds—a particularly strong method like that of Wilkinson (2021) will be able to compare many of them, while a weak method will not. But even with a weak method, such as that given by Additivity, we can begin to offer plausible verdicts in various difficult cases described above. And so we may be able to restore the judgements of aggregative theories in practice.

7 Appendix

Theorem 2: For any reflexive, transitive relation \succsim on \mathcal{P} , if \succsim satisfies Stochastic Dominance, Finite Expectations, and Separability of Value (for Lotteries) then it satisfies ED1.

Proof: Let $L_1, L_2 \in \mathcal{P}$ be any lotteries for which:

$$\sum_{W_j \in \mathcal{W}_1} L_1(W_j) \left(\sum_{l \in \mathcal{L}} V_j(l) - V_*(l) \right) \geq \sum_{W_j \in \mathcal{W}_2} L_2(W_j) \left(\sum_{l \in \mathcal{L}} V_j(l) - V_*(l) \right) \quad (\text{i})$$

Either 1) both sides of the inequality converge unconditionally to some real values, with LHS \geq RHS, or 2) LHS diverges unconditionally to $+\infty$, RHS diverges unconditionally to $-\infty$, or both.

If (1): define $L_1^*, L_2^* \in \mathcal{P}$ by $L_1^*(W - W_*) = L_1(W)$ and $L_2^*(W - W_*) = L_2(W)$ for all $W \in \mathcal{W}_{(1,2)}$ (and $L_1^*(W), L_2^*(W) = 0$ otherwise). Since both sides of (i) converge unconditionally to real values, the expected total sums of local value in L_1^* and L_2^* will be finite, and the expected total sum for L_1^* will be greater than or equal to that for L_2^* . By Finite Expectations, $L_1^* \succsim L_2^*$. Then, by Separability of Value (for Lotteries), $L_1 \succsim L_2$, as required.

If 2): LHS of (i) diverges to $+\infty$, or the RHS diverges to $-\infty$, or both. If both, then define L_1^* as above and $L_2^*(W_0) = 1$. If not both, then define L_1^*, L_2^* as above. Then $L_1^* \succsim L_2^*$.

If LHS diverges to $+\infty$ then, by Stochastic Dominance there is some lottery L_1^{**} with finite expected sum such that $L_1^* \succsim L_1^{**}$ and $L_1^{**} \succsim L_2^*$ by Finite Expectations. We can obtain L_1^{**} from L_1^* by replacing any worlds in its domain that have total value greater than a given finite bound with some other worlds with totals below that bound.

If LHS does not diverge but RHS does diverge to $-\infty$, by Stochastic Dominance there is some lottery L_2^{**} such that $L_2^* \preccurlyeq L_2^{**}$ and $L_2^{**} \preccurlyeq L_1^*$ by Finite Expectations.

Either way, $L_1^* \succsim L_2^*$. Separability of Value (for Lotteries) then implies that $L_1 \succsim L_2$. \square

* * *

The proof of Theorem 4 below relies on Lemma 1. This lemma implies that, if Independence and Extrapolated Expectations hold, then we can evaluate any lottery in which every outcome is some scalar multiple of some world $W \succ \mathbf{0}$. We evaluate it as the probability-weighted sum of scalar multiples of W .

Lemma 1: Let L be any lottery in

$$\mathcal{P}$$

with finite domain and such that, for some $W' \in \mathcal{W}$, $L(W) > 0$ only if $W = k \cdot W'$ for some positive, real k_j . If \succsim satisfies Independence and Extrapolated Expectations, then:

$$L \sim \left(\sum_{W_j \in \mathcal{W}} L(W_j) \times k_j \right) \cdot W$$

Proof:

Let k_{\max} be the greatest such k_j (which exists, since L has finite domain).

For each k_j such that $L(k_j \cdot W') > 0$, let $L_j(k_{\max}) = \frac{k_j}{k_{\max}}$ and $L_j(\mathbf{0}) = 1 - \frac{k_j}{k_{\max}}$. Since $k_j \cdot W = \frac{k_j}{k_{\max}} \cdot (k_{\max} \cdot W)$, Extrapolated Expectations implies that $k_j \cdot W \sim L_j$.

By Independence, in L we can replace each world $k_j \cdot W$ with lottery L_j . In other words, by (i) and Independence, $L \sim L'$, where $L'(k_{\max} \cdot W') = \sum_{k_j \in \mathbb{R}} L(k_j \cdot W') \times \frac{k_j}{k_{\max}}$ and $L'(\mathbf{0}) = 1 - L'(k_{\max} \cdot W')$.

By Extrapolated Expectations, $L' \sim \left(\sum_{k_j \in \mathbb{R}} L(k_j \cdot W) \times \frac{k_j}{k_{\max}} \right) \cdot (k_{\max} \cdot W)$.

$\therefore L \sim \left(\sum_{W_j \in \mathcal{W}} L(W_j) \times k_j \right) \cdot W$, as required. \square

Theorem 4: For any reflexive, transitive relation \succsim on \mathcal{P} , if \succsim satisfies Stochastic Dominance, Separability of Value, Independence, and Extrapolated Expectations, then it satisfies ED2.

Proof:

For any $L_1, L_2 \in \mathcal{P}$ and any $W_i \in (\infty, \in)$, define $\Delta p_i = L_1(\succ W_i) - L_2(\succ W_i)$. And let $L_1, L_2 \in \mathcal{P}$ be lotteries such that, for some W_* and any $W_i \in (\infty, \in)$,

$$k_i \Delta p_i \cdot D_* \preceq \Delta p_i \cdot D_i \quad (\text{i}) \quad \text{and} \quad \sum_{W_i \in \mathcal{W}_{(1,2)}} k_i \Delta p_i \geq 0 \quad (\text{ii})$$

For some such W_* and k_i s, let $L_1^D \left(\left(\sum_{W_j \in \preceq W_i} k_j \right) \cdot D_* \right) = L_1(W_i)$ for all $W_i \in \mathcal{W}_1$, and similarly for L_2^D .

By Separability of Value, $D_* \succ \mathbf{0}$ (since it is the difference between some world and a worse one).

So, by Lemma 1,

$$L_1^D \sim \left(\sum_{W_i \in \mathcal{W}_1} L_1(W_i) \left(\sum_{\{W_j \in \mathcal{W}_{(1,2)}\}} k_j \right) \right) \cdot D_* \quad \text{and similarly for } L_2^D \text{ and } L_2.$$

So $L_1^D \succcurlyeq L_2^D$ if and only if:

$$\begin{aligned} & \sum_{W_i \in \mathcal{W}_1} L_1(W_i) \left(\sum_{\{W_j \in \mathcal{W}_{(1,2)}\}} k_j \right) \geq \sum_{W_i \in \mathcal{W}_2} L_2(W_i) \left(\sum_{\{W_j \in \mathcal{W}_{(1,2)}\}} k_j \right) \\ \Leftrightarrow & k_{\min} + \sum_{W_i \in \mathcal{W}_{(1,2)}} L_1(\succcurlyeq W_i) \times k_i \geq k_{\min} + \sum_{W_i \in \mathcal{W}_{(1,2)}} L_2(\succcurlyeq W_i) \times k_i \quad \text{where } k_{\min} = \min\{k_i | W_i \in \mathcal{W}_{(1,2)}\} \\ & \Leftrightarrow \sum_{W_i \in \mathcal{W}_{(1,2)}} k_i \Delta p_i \geq 0 \end{aligned}$$

Therefore, by (ii), $L_1^D \succcurlyeq L_2^D$.

For the same W_* and k_i s, define $L_1^*(W - W_*) = L_1(W)$ for all $W \in \mathcal{W}_1$, and L_2^* likewise (mutatis mutandis). These lotteries resemble L_1 and L_2 ; each outcome just has W_* is subtracted from it. As a result, each outcome $W - W_*$ can be represented as:

$$W - W_* = \sum_{\{W_j \in \mathcal{W}_{(1,2)} | W_* \preccurlyeq W_j \preccurlyeq W\}} D_j \quad (\text{or, for } W \prec W_*, \text{ as } \sum_{\{W_j \in \mathcal{W}_{(1,2)} | W \preccurlyeq W_j \preccurlyeq W_*\}} -D_j).$$

But, by (i), whenever $\Delta p_i > 0$,

$$W_i - W_* = \sum_{\{W_j \in \mathcal{W}_{(1,2)} | W_* \preccurlyeq W_j \preccurlyeq W_i\}} D_j \succcurlyeq \left(\sum_{\{W_j \in \mathcal{W}_{(1,2)} | W_j \preccurlyeq W_i\}} k_j \right) \cdot D_*$$

and, whenever $\Delta p_i < 0$, the inequality is reversed. Therefore, by Stochastic Dominance and (i), $L_1^* \succcurlyeq L_1^D$ and $L_2^* \succcurlyeq L_2^D$. Then, by transitivity of \succcurlyeq , $L_1^* \succcurlyeq L_2^*$.

By Separability of Value, $L_1 \succcurlyeq L_2$ if and only if $L_1^* \succcurlyeq L_2^*$. Therefore, $L_1 \succcurlyeq L_2$, as required. \square

References

- ARNTZENIUS, F., 2014. Utilitarianism, decision theory and eternity. *Philosophical Perspectives*, 28, 1 (2014), pp. 31–58. (cited on pages 2, 3, 6, 7, 8, and 23)
- ASKELL, A., 2019. *Pareto Principles in Infinite Ethics*. PhD dissertation, New York University. (cited on page 4)
- BASU, K. AND MITRA, T., 2007. Utilitarianism for infinite utility streams: A new welfare criterion and its axiomatic characterization. *Journal of Economic Theory*, 133, 1 (2007), pp. 350–73. (cited on page 5)

- BLACKORBY, C.; BOSSERT, W.; AND DONALDSON, D., 1995. Intertemporal population ethics: Critical-level utilitarian principles. *Econometrica*, 63, 6 (1995), pp. 1303–20. (cited on page 1)
- BOSTROM, N., 2011. Infinite ethics. *Analysis and Metaphysics*, 10 (2011), pp. 9–59. (cited on pages 2, 3, 6, 7, 8, and 23)
- BROOME, J., 2004. *Weighing Lives*. Oxford University Press, Oxford. (cited on page 4)
- BUCHAK, L., 2013. *Risk and Rationality*. Oxford University Press, Oxford. (cited on page 13)
- CARROLL, S., 2017. Why Boltzmann brains are bad. *arXiv preprint*, (2017). (cited on page 2)
- DE SIMONE, A.; GUTH, A.; LINDE, A.; NOORBALA, M.; SALEM, M.; AND VILENKIN, A., 2010. Boltzmann brains and the scale-factor cutoff measure of the multiverse. *Physical Review D*, 82, 6 (2010), 063520. (cited on page 2)
- DREIER, J., 2019. World-centered value. In *Consequentialism: New Directions and Problems* (Ed. C. SEIDEL), p. 31–50. Oxford University Press, Oxford. (cited on page 4)
- GARRIGA, J. AND VILENKIN, A., 2001. Many worlds in one. *Physical Review D*, 64, 4 (2001), 043511. (cited on page 2)
- GUTH, A., 2007. Eternal inflation and its implications. *Journal of Physics A: Mathematical and Theoretical*, 40, 25 (2007), pp. 6811. (cited on page 2)
- HUEMER, M., 2008. In defence of repugnance. *Mind*, 117, 468 (2008), pp. 899–933. (cited on page 4)
- JACKSON, F., 1991. Decision-theoretic consequentialism and the nearest and dearest objection. *Ethics*, 101, 3 (1991), pp. 461–82. (cited on page 3)
- JONSSON, A. AND VOORNEVELD, M., 2018. The limit of discounted utilitarianism. *Theoretical Economics*, 13, 1 (2018), pp. 19–37. (cited on pages 2, 6, and 23)
- LAUWERS, L. AND VALLENTYNE, P., 2004. Infinite utilitarianism: more is always better. *Economics & Philosophy*, 20, 2 (2004), pp. 307–30. (cited on pages 2, 5, and 23)
- LAUWERS, L. AND VALLENTYNE, P., 2016. Decision theory without finite standard expected value. *Economics and Philosophy*, 32, 3 (2016), pp. 383–407. (cited on page 13)
- MACASKILL, W.; VALLINDER, A.; SHULMAN, C.; ÖSTERHELD, C.; AND TREUTLEIN, J., n.d. The evidentialist’s wager. Unpublished manuscript, available at https://globalprioritiesinstitute.org/wp-content/uploads/2019/MacAskill_et_al_Evidentialist_Wager.pdf. (cited on page 2)
- MEACHAM, C., 2020. Too much of a good thing: Decision-making in cases with infinitely many utility contributions. *Synthese*, (2020). (cited on pages 3 and 7)
- NEBEL, J., 2018. The good, the bad, and the transitivity of *better than*. *Noûs*, 52, 4 (2018), pp. 874–99. (cited on page 4)
- PARFIT, D., 1997. Equality and priority. *Ratio*, 10 (1997), pp. 202–21. (cited on page 1)
- PIVATO, M., 2014. Additive representation of separable preferences over infinite products. *Theory and Decision*, 77, 1 (2014), pp. 31–83. (cited on page 8)

- SCHOENFIELD, M., 2014. Decision making in the face of parity. *Philosophical Perspectives*, 28 (2014). (cited on pages 12 and 13)
- SEIDENFELD, T.; SCHERVISH, M.; AND KADANE, J., 2009. Preference for equivalent random variables: A price for unbounded utilities. *Journal of Mathematical Economics*, 45 (2009), pp. 5–6. (cited on page 13)
- SMITH, N., 2014. Is evaluative compositionality a requirement of rationality? *Mind*, 123, 490 (2014), pp. 457–502. (cited on page 13)
- TEMKIN, L., 2014. *Rethinking the Good*. Oxford University Press, Oxford. (cited on page 4)
- VALLENTYNE, P. AND KAGAN, S., 1997. Infinite value and finitely additive value theory. *Journal of Philosophy*, 94, 1 (1997), pp. 5–26. (cited on pages 2, 5, 6, and 23)
- WALD, R., 1983. Asymptotic behavior of homogeneous cosmological models in the presence of a positive cosmological constant. *Physical Review D*, 28, 8 (1983), pp. 2118. (cited on page 2)
- WILKINSON, H., 2021. Infinite aggregation: Expanded addition. *Philosophical Studies*, 178, 6 (2021), pp. 1917–49. (cited on pages 2, 4, 6, 23, and 24)