



Infinite aggregation: expanded addition

Hayden Wilkinson¹ 

Published online: 25 August 2020
© Springer Nature B.V. 2020

Abstract How might we extend aggregative moral theories to compare infinite worlds? In particular, how might we extend them to compare worlds with infinite spatial volume, infinite temporal duration, and infinitely many morally valuable phenomena? When doing so, we face various impossibility results from the existing literature. For instance, the view we adopt can endorse the claim that (1) worlds are made better if we increase the value in every region of space and time, or (2) that they are made better if we increase the value obtained by every person. But they cannot endorse both claims, so we must choose. In this paper I show that, if we choose the latter, our view will face serious problems such as generating incomparability in many realistic cases. Opting instead to endorse the first claim, I articulate and defend a spatiotemporal, expansionist view of infinite aggregation. Spatiotemporal views such as this do face some difficulties, but I show that these can be overcome. With modification, they can provide plausible comparisons in cases that we care about most.

Keywords Infinite utility streams · Utilitarianism · Intergenerational equity · Pareto · Aggregation · Infinite value

1 Introduction

Suppose you found that the universe around you was infinite—that it extended infinitely far in space or in time and, as a result, contained infinitely many persons. How should this change your moral decision-making? Radically, it seems, if you

✉ Hayden Wilkinson
hayden.wilkinson@anu.edu.au

¹ School of Philosophy, Australian National University, Canberra, ACT, Australia

accept a moral theory which gives some consideration to the total value in the world.

Let's assume that our universe is that way. Further, assume that infinitely many of the persons within it will have quite good lives (with positive value greater than some fixed $\epsilon > 0$). The total sum of value in the world is then positively infinite (or else undefined). But suppose you make any change you want to the world. If infinitely many positive-valued lives will still exist, then the total value will still be positively infinite (or else undefined). So you cannot compare the two worlds by their totals. Neither contains greater total value.

Here is a more concrete case. You can either rescue one person from death, or rescue five others from death (each of whom would be better off not dying). And, either way, infinitely many other persons with valuable lives will exist throughout the universe. We can represent the outcome of saving one with W_1 , and that of saving five with W_5 . Each world contains an infinite plurality of persons $\{p_a, p_b, p_c, \dots\}$. And here the moral value of each person's life is represented on an interval scale as 0 (if they die now) or 1 (if they get to enjoy the remainder of their life).

	p_a	p_b	p_c	p_d	p_e	p_f	p_g	p_h	p_i	p_j	\dots
W_1 :	1	0	0	0	0	0	1	1	1	1	\dots
W_5 :	0	1	1	1	1	1	1	1	1	1	\dots

For these worlds, we cannot say that either has the greater total sum. Even if we extend the real numbers by defining ∞ (or even by defining the transfinite ordinals or cardinals as well), the total sum of these two worlds will still be the same infinite number; their totals will be equal. If our only means of evaluating worlds is by their total, then we cannot say which of these worlds is better. But W_5 is clearly better! Five people are better off, and only one worse off (by the same amount). Intuition says that W_5 is better and that, absent any other considerations, we ought to save the five.

You will often have such difficulty in infinite worlds if the true moral theory relies on total value, at least in part, to judge actions. The problem is clearest for maximizing, aggregative, pure consequentialists, who consider the total sum of value to be the sole determinant of moral betterness (and betterness the sole determinant of what we ought to do).

But such consequentialists are not the only victims. A much broader class of moral theories—call them *minimally aggregative* views—recognize that we have a *pro tanto* reason to impartially promote value.¹ Those theories may recognize other considerations—constraints, prerogatives, satisfactory thresholds of value, and reasons stemming from other sources. But they still recognize that, when other considerations are silent, we ought to bring about one world rather than another if and only if it contains greater total value. This holds most plausibly in pure rescue cases, such as above. (Assume that no rights are violated, no duties unfulfilled, no

¹ We need not agree on what composes value—it may be pleasure, happiness, a plurality of things, or any other quantity associated with all tokens of some physical phenomena.

special relationships present, etc.) In such a case, many impurely consequentialist and non-consequentialist views still say that you ought to save the five, and that you ought to *because* W_5 is a better outcome—it contains more of the good. But, if the world is infinite in the relevant way, it doesn't. If so, even *minimally* aggregative views fail to give plausible judgments in cases like this. And that seems a compelling *reductio* of all such views.

If you hold such a view, you might be disappointed to find that the world around you is infinite in the relevant sense. I am sorry to disappoint you, but contemporary physics suggests just that. The widely accepted *flat-lambda* model predicts that our universe will tend towards a stable state and will then remain in that state for infinite duration (Wald 1983; Carroll 2017). Also widely accepted, the *inflationary view* posits that our world is spatially infinite, containing infinitely many other 'bubble' universes beyond our cosmic horizon (Guth 2007). But that's not all they predict. Take any small-scale phenomenon which is morally valuable e.g., perhaps a human brain experiencing the thrill of reading philosophy for a given duration. Each of the above physical views predicts that our universe, in its infinite volume, will contain infinitely many² such thrills (Garriga and Vilenkin 2001; Linde 2007; de Simone 2010; Carroll 2017).³ If we sum up the value of all of those thrills, we have an infinite total. So, our world will contain infinite total value, at least by those physical theories. But those theories are widely accepted among physicists, so this seems likely to be true. And if it is, many moral theories fail us.

There are views not too distant from minimally aggregative views which might avoid the infinitary problem—for instance, views which do not aggregate all value impartially. You might *discount* value far in the distance or the future when you aggregate, via a non-zero rate of pure time preference (e.g., Koopmans 1972; see Section 3). Or you might exclude value from lives which haven't yet begun or which exist only in the outcomes you won't actually choose, or whose existence depends on your actions⁴ (e.g., Heyd 2009; Bader 2020). That is, you might adopt one of the many person-affecting views of population ethics, and accept the controversial implications it brings (see Beckstead 2013).

But this may not be necessary. We might be able to retain aggregative views, minimal or otherwise, in their fully impartial form. And if not, then we may have a compelling argument for discounting or person-affecting views.

We might be able to do this by revising our method of comparing worlds to not rely on real-valued representations of total value.⁵ We have various proposals for

² The set of such token events will be *countably* infinite. For any given phenomenon, I'll assume, each token must occupy some (exclusive) non-zero volume of spacetime. For illustration, a human brain to experience a given 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.

³ Physicists are hesitant to speak of infinite quantities (exceptions include Garriga and Vilenkin 2001; Gott 2008), as these can never be empirically confirmed. Nonetheless, this conclusion is entailed by both of the above views.

⁴ The latter two options won't help with the idealized example considered above, but they do help with cases in which we change the identities of many future persons. See Sect. 5.3.

⁵ Following others, I pursue an *axiological* solution rather than a *normative* one—one which revises how we compare worlds, rather than directly revise how we judge the permissibility of actions (whether

how to do this (e.g., Vallentyne 1993; Vallentyne and Kagan 1997; van Liedekerke and Lauwers 1997; Lauwers and Vallentyne 2004; Bostrom 2011; Arntzenius 2014; Jonsson and Voorneveld 2018; see Askell 2018 for a survey). I won't explore all of these approaches here. Instead, I focus on one promising approach: *expansionism* (cf. Vallentyne and Kagan 1997; Arntzenius 2014).

Specifically, I propose and defend a *spatiotemporal* version of expansionism.⁶ This contrasts with the previous version given by Vallentyne and Kagan, which is agnostic as to what the basic locations of value are. Those locations might be individual persons, or positions in spacetime, or something else. Without knowing which, we often cannot apply the rule in practice. And if we adopt the obvious choice of locations—persons—then it turns out that we cannot give plausible judgments (or any judgments at all) in even some fairly mundane cases (see Sect. 3 below). So I argue for a *spatiotemporal* version of expansionism—one which adopts spacetime positions as our basic locations, and uses their natural structure to give more plausible judgments than we otherwise could.

Approaches like mine face objections (see Sect. 3). One is that they seem to abandon a crucial feature of aggregation—its *impartiality*, including over properties such as *where* each person happens to be. Another is that they give counterintuitive judgments in key cases (see Cain 1995). And another objection which I introduce in this paper is that, in some fairly mundane cases we might actually face, my approach seems to fall silent. I'll argue that these objections are not decisive. For one, we will see that basic tenets like impartiality do not rule out spatiotemporal views entirely. For two, we must accept counterintuitive judgments in some cases to avoid far more implausible judgments in other (more realistic) cases. And, for three, with a bit of work we can construct at least some spatiotemporal views which *do* give plausible judgments even in many problem cases.

I hope to thereby show that: we can solve the basic problem of infinite aggregation; we can restore the judgment that it is better to save five than to save one; and we can do so without many of the problems of other proposals, such as delivering incomparability in almost all cases we ever face (see Sects. 3, 6). In our actual universe and in cases we may actually face, we can still make decisions based on what will promote the good.

Footnote 5 continued

objectively or subjectively). This is for three reasons. The first: brevity. The second: for all minimally aggregative views, the problem arises from having an axiology which doesn't discriminate enough; with an axiological solution, we'll automatically have a normative solution. And the third: these views are united by their axiology, but differ in all else, so altering the axiology will provide the most general solution.

There is one proposed solution which directly modifies our (subjective) normative principle rather than axiology, given by Arntzenius (2014). I argue elsewhere that this solution is untenable (Wilkinson n.d.(c)).

⁶ When discussing expansionism, Askell (2018, pp. 71–87) also focuses on the spatiotemporal version. I agree that this is the most compelling version, as I argue in Sect. 3 and Wilkinson (n.d.(a)).

2 Preliminaries

2.1 Locations

We want to compare worlds on the basis of the value they contain. We cannot compare their totals, so we look to the individual instances of value—*local values*. We could individuate and identify these across worlds in different ways, e.g.: by who the *person* is whose life contains the value; or which person-time-slice obtains it; or in which generation⁷ it is obtained; or at which *position* in space and/or time they arise, or perhaps some other way. For any of these, local value is associated with token entities of some common type which exist (or have counterparts) across different worlds. I'll call those tokens *locations*. For reasons given in Sect. 3, I focus on the two most plausible types of locations: persons, and spacetime positions.

Whichever locations we adopt, each world W_i contains a set of them: $\mathcal{L}_i = \{l_1, l_2, l_3, \dots\}$.⁸ The subscripts j for each l_j may be arbitrary or may reflect some natural, essential structure of locations (as for spacetime positions). The worlds we compare will often contain the same (or counterpart) locations, in which case we can use just the one set \mathcal{L} . (More on this below.)

For each world W_i , some value function $V_i : \mathcal{L}_i \rightarrow \mathbb{R}$ maps each location to its value. So, in W_1 , the local value at location l_a is given by $V_1(l_a)$. I assume that local values admit at least a cardinal representation on the reals—that is, a representation which is unique at least up to affine transformations. This means that, for any given location l_a which appears in each such world, either $V_1(l_a) - V_2(l_a) \geq V_3(l_a) - V_3(l_a)$ or vice-versa. But I will not assume that local values have a natural zero or a natural unit (as is necessary for ratio-scale or translation-scale representations, respectively).

I also assume that locations—whatever type they might be—can be positioned in space and time within each world. If locations themselves are spacetime positions, that's easy enough—they each have an essential position. If instead they are persons, they still each occupy spacetime positions. For simplicity, let's treat each person's position as a single point: perhaps the point of their birth, or the midpoint of all positions they ever occupy. (Nothing below hangs on which we choose.) So, whichever form our locations take, they are assigned a point in spacetime $\mathbf{x} \in \mathbb{R}^4$, or (x, y, z, t) . This is represented in a Cartesian coordinate system on four dimensions and is unique up to translation and scalar multiplication.⁹

⁷ Although generations are commonly used as locations in the social welfare literature, I agree with Askill (2018, pp. 12–13) that they are typically underspecified and, when precisified, bring implausible results.

⁸ This set may be countable or uncountable. If the latter, we won't be able to list our locations with natural-numbered subscripts. But much of what follows does not depend on which it is.

If our locations are persons, we can assume that the set will be countable, for reasons similar to those raised in Footnote 2: persons require some positive, finite volume and duration. We could set some minimum physical volume and duration for a person to matter morally, and only a countable infinity of persons will fit into any spacetime (whether Euclidean or not).

If our locations are spacetime positions, we can assume that the set will be uncountable. After all, these are positioned densely.

⁹ I am assuming that spacetime is *Euclidean*: that points have a unique coordinate representation up to translation and scalar multiplication; and that every two points have a unique inner product and geometric

2.2 Identity/counterpart relations

My strategy for comparing worlds—the strategy of almost everyone in the literature to date—is to take sets of individual locations and to compare their local values across worlds. But, to compare local values across worlds, we need the same locations (or, alternatively, their *counterparts*) to exist across worlds. How do we match up locations with each other?¹⁰ If we cannot answer this, then my strategy will fail immediately. I'll describe what I think are the most plausible methods of identifying locations with themselves or their counterparts, whether the relevant type of locations is persons or spacetime positions.

Suppose that locations are individual *persons*. In this case, I propose that we stick to our theory of personal identity, whichever is correct—that our identity or counterpart relation use criteria analogous to those used by the identity relation between a person's past, present, and future selves. If all that's required for one's past and future selves being the same person is that they share some bundle of qualitative properties, let that be sufficient for a person in one world to be the same as a person in another world. Alternatively, the identity relation between past and future selves might require physical continuity. If so, then let our transworld relation hold between two persons if and only if they share some common history of identical events. (To determine which events are identical, we can use the following proposal.)

Alternatively, suppose that our locations are *spacetime positions*. Then we may not have such a theory of identity. But that is no trouble—there is an obvious identity/counterpart relation, at least in any pairs of worlds we'll ever need to compare. Since our actions necessarily cannot change the past, any such worlds will share the same past events (at the same positions). In worlds like these with common histories up to the present, let us map all past positions to those occupied by the same events (e.g., we can map the position of Runnymede in 1215 in one world to the same position in every other world, as that's where the signing of the Magna Carta occurs in all of them). Then we can map future positions too: each such x is uniquely specified by its spatial and temporal distance from (any four) past points,¹¹ so we can specify its transworld identities/counterparts as the positions which are also those same distances from those same points.

Footnote 9 continued

distance between them (up to scalar multiplication). This distinguishes the spacetime I'm working with from Minkowskian and pseudo-Riemannian spacetimes which are more accurate representations of our universe and which each admit different coordinate representations under which inner products and geometric distance are not fixed. I accept that the assumption of Euclidean spacetime is unrealistic, but necessary to avoid taking us too far afield. I wish to demonstrate that my method overcomes basic problems in even the Euclidean setting before attempting to deal with problems introduced by relativistic spacetimes. In further work, I show that these further problems can be resolved as well (Wilkinson n.d.(b)).

¹⁰ Note that I am not committing myself to a particular position in the debate—between Lewis (1971) and, for instance, Kripke (1980, p. 45)—on whether transworld identities exist. The following relation works just as well whether it's an identity relation or just a privileged counterpart relation.

¹¹ This clearly holds in Euclidean spacetime. In relativistic spacetimes, the counterpart relation becomes more complicated (see Wilkinson n.d.(b)).

This matches our intuitive sense of the ‘same’ positions (e.g., “1 meter in front of me, 1 s in the future” will be the same point no matter what I do). It also preserves the distances among all spacetime points. And it ensures that any two worlds we must ever compare—the outcomes of two actions available to us—will share the same set of positions, which will prove useful in what follows. The same cannot be said for worlds which do not have some shared history, but in practice we’ll never need to choose among such worlds.¹²

2.3 Desiderata

To compare worlds, we want an ‘at least as good as’ relation \succsim on the set of (metaphysically possible) worlds \mathcal{W} . I assume that this relation is reflexive and transitive.¹³ The asymmetric component (\succ) will be our strict betterness relation, and the symmetric component (\simeq) our equality relation.

Defining a satisfactory betterness relation is no easy feat. We must make some hard choices. From the social welfare literature, we have a collection of nasty impossibility results—in particular, from Zame (2007) and Lauwers (2010).¹⁴ They show that it is impossible to construct¹⁵ a version of \succsim which is transitive, complete, and does not violate either *Pareto* or *Finite Anonymity*—both of which are basic requirements for a plausible betterness relation.¹⁶

Pareto holds that, if two worlds contain precisely the same locations all with precisely the same local values, then those worlds are equally good; and if we made one of those worlds better at some locations, then it would then better than the other. And this seems highly plausible, for any type of location! It can be stated more formally as follows.

Pareto (over locations): For any worlds W_1 and W_2 containing the same locations \mathcal{L} , if $V_1(l) \geq V_2(l)$ for all $l \in \mathcal{L}$, then $W_1 \succsim W_2$.

If, as well, $V_1(l_i) > V_2(l_i)$ for some $l_i \in \mathcal{L}$, then $W_1 \succ W_2$.

Meanwhile, *Finite Anonymity* is necessary for our aggregation to be genuinely impartial. If two worlds only differ by swapping the local values of two locations, then those worlds must be equally good.

¹² Note also that this proposal does not require that counterpart positions are metaphysically identical, nor that they have essential properties. I propose merely that, in our moral evaluations, we treat points as equivalent when they satisfy these relations. This is compatible with both spacetime substantivalism and relationism.

¹³ Transitivity has its critics, e.g., Temkin (2012). I nonetheless find it overwhelmingly plausible. In keeping with the infinite aggregation literature to date (within both moral philosophy and welfare economics), I will assume without argument that it holds.

¹⁴ Lauwers shows that the following impossibility holds even for the domain of worlds with local values of only 0 and 1.

¹⁵ It has been shown that such relations do exist, but cannot be constructed (see Svensson 1980). This is unsatisfactory for an moral rule, by which we would like to explicitly define the criteria for betterness.

¹⁶ These results are originally given in a setting in which we apply Pareto and Finite Anonymity over generations. But as Askell (2018, pp. 39–44) explains, with some minimal assumptions, we obtain the same result if we take persons as locations. Likewise, it applies if spacetime positions are locations.

Finite Anonymity (over locations): If for some $l_a, l_b \in \mathcal{L}$ we have $V_1(l_a) = V_2(l_b)$, $V_1(l_b) = V_2(l_a)$ and $V_1(l) = V_2(l)$ for all other $l \in \mathcal{L}$, then $W_1 \simeq W_2$.

Since \succsim must be transitive, Finite Anonymity implies that we can permute *any* finite number of local values and the resulting world remains equally good. But note that it does not imply that the result of an *infinite* number of permutations—shuffling infinitely many local values—will be equally good. (More on this below.)

Like Pareto, Finite Anonymity seems highly plausible. I find both Pareto and Finite Anonymity hard to deny, at least for *some* type of location (if not multiple types). Perhaps they hold for persons, or perhaps for positions. (Note that, in finite worlds, they hold for *all* types.) Whichever it is, the Zame–Lauwers impossibility will apply. Does this mean that any proposal for a betterness relation over infinite worlds is doomed from the start? I don't think so. To avoid the impossibility while hanging onto these principles, we can simply allow that \succsim is *not complete*. This would mean that there is some W_1 and W_2 in \mathcal{W} such that neither $W_1 \succsim W_2$ nor $W_1 \precsim W_2$. This may seem counterintuitive, but it is better than the alternative: having \succsim violate either Pareto or Finite Anonymity (for every type of location), or else fail to be transitive. So, throughout this paper, I won't be seeking a complete relation. No such relation would be plausible.

But a moral theory still needs to give verdicts in mundane cases that agents like us *actually* face. Even if we cannot demand that the \succsim relation be *complete* (over the set of all metaphysically possible worlds), it must still be '*minimally complete*'. By this I don't mean anything as simple as completeness over the set of all epistemically possible worlds, or all physically possible worlds. Instead, I mean this: for any decision-making agent who will *actually* exist, for any decision they will actually face, and for any two worlds which would be brought about if they took one or another action in that decision, those worlds must be comparable by the \succsim relation. If so, then \succsim is minimally complete. Or, in short: \succsim must never *actually* fall silent on us. It's fine for it to be silent in exotic problem cases—such as when deciding between worlds which don't share the same history, or between highly gerrymandered worlds. But, if it falls silent in decisions I actually need to make, that's a serious failing.

Note that it would be impossible to prove that any given \succsim relation is minimally complete. Instead, I will simply show that some proposed \succsim relations (in the next section and Sect. 5) fail to be.

To sum up, my desiderata for infinite aggregation are that:

- the betterness relation is reflexive and transitive;
- it appears minimally complete, and gives plausible judgments in those cases it covers;
- it satisfies Pareto *over spacetime positions*; and
- it satisfies Finite Anonymity *over spacetime positions*.

Why over spacetime positions but not (also) over persons? And why Finite Anonymity rather than a stronger, infinite form of anonymity? I'll justify these choices below.

3 Isn't spatiotemporal position morally irrelevant?

My view gives verdicts which are sensitive to where value is positioned in spacetime, as do almost all views proposed in the social welfare literature. In this section, I will pre-emptively (and all too briefly) address the most common objections to such views. For an extended defence, see Wilkinson (n.d.(a)).

3.1 Pareto: don't take it personally

We cannot endorse both Pareto *over spacetime positions* and Pareto *over persons*. The two principles are incompatible. (Likewise, Pareto over any location-type will be incompatible with Pareto over almost any other.) But Pareto over persons seems highly plausible, so this constitutes a serious objection to my view and, indeed, to many views in the literature.

To see why the two principles are incompatible, consider *The Shuffle*.¹⁷

Example: The Shuffle

You are hosting a line dancing event, to which infinitely many people have shown up—many more than you expected! Fortunately, you have two dances prepared which can accommodate infinitely many dancers. Each dance involves the dancers standing in single file. And, for both, the moves required of each dancer depend on their position in the line, so dancers will find some positions more strenuous and less pleasant than others. If you ask the dancers to perform one dance or the other then the moral value obtained by the dancers (and positions) is given by W_1 or W_2 , respectively.

	p_1	p_2	p_3	p_4	p_5	p_6	p_7	p_8	p_9	p_{10}	\dots
	x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8	x_9	x_{10}	\dots
$W_1 :$	1	0	1	0	1	0	1	0	1	0	\dots
	p_1	p_2	p_4	p_6	p_3	p_8	p_{10}	p_{12}	p_5	p_{14}	\dots
	x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8	x_9	x_{10}	\dots
$W_2 :$	1	0	0	0	1	0	0	0	1	0	\dots

Here, the worlds W_1 and W_2 contain precisely the same persons and the same spatiotemporal positions. But many of those persons are at different positions in each world—they're *shuffled*. Compared to W_1 , every even-numbered person from p_4 onwards is positioned earlier in the line. And every odd-numbered person from p_3 onwards is positioned later.

¹⁷ Cases like this, which demonstrate the disagreements between different Pareto principles, originate with Cain (1995, pp. 401–403). Cain considers them a decisive objection to spatiotemporal views like mine. But, as we'll see below, views that endorse Pareto over persons face problems that you might consider even worse.

Which dance would bring the better outcome? Look at the spacetime points: all points do at least as well in W_1 , and some do strictly better ($\mathbf{x}_3, \mathbf{x}_7, \dots, \mathbf{x}_{4n+3}, \dots$). Pareto *over positions* says that $W_1 \succ W_2$.

Look at the persons: every person bears the same amount of value in W_1 as in W_2 , they're just positioned less densely. Pareto *over persons* says that $W_2 \simeq W_1$, in conflict with Pareto over positions.

The judgment from Pareto over persons may seem compelling—no person would be worse off in W_2 , so what's the harm in choosing that dance? As a person myself, I wouldn't mind finding myself in W_2 . And you might think that improving outcomes for people is what matters morally, rather than simply increasing value in the abstract. You might think, as Askell (2018, p. 84) does, that "...Ethics is concerned with people. It is not particularly concerned with the pattern of utility across spacetime." I have sympathy with this thought, but we also have reason to reject it. I provide only a brief argument here, but see Wilkinson (n.d.(a)) for a full defence.

The reason is this: unlike Pareto over positions, Pareto *over persons* implies *radically* widespread incomparability among worlds, even in quite mundane cases. So any theory that satisfies Pareto over persons fails to be minimally complete.

To argue this, I must first make two assumptions explicit. The first: \succ is a *qualitative* relation. It ranks any pair of worlds (W_1, W_2) the same way as it would rank any other world pair (W_3, W_4) with the same qualitative properties. Note that it need not also be a *qualitative internal relation*, which means that it must rank worlds W_1 and W_2 as equally good if they are qualitative duplicates. No, that would violate Pareto immediately. Instead I am assuming merely that our rankings of worlds will not change if we change the specific identities of persons within them in the *same way* in *both* worlds.¹⁸

The second: we can change many of the qualitative features of persons however we want (including at least their positions) without changing their identity (or counterpart relationships). This implies that worlds W_1 and W_2 above are both metaphysically possible—we can shift infinitely many of the even-numbered persons leftwards by any distance, and odd-numbered persons rightwards, while preserving their identities. Both this and the previous assumption are hard to deny, particularly if we think that such problem cases are possible.

Now consider *Firing Line*.

Example: Firing Line

Infinitely many persons are lined up in front of you, extending infinitely far into the distance. You have a choice to make. Either let every *second* person in the line be killed, or let every *fifth* person be killed. Their lives will otherwise be quite pleasant, and equally so.

Here are the local values for *Firing Line*. The i th person in the line corresponds to p_i , at position \mathbf{x}_i . And note that we have the beneficiaries in W_2 are more *densely*

¹⁸ This assumption is often called *Relative Anonymity* in the social welfare literature, e.g., by Asheim et al. (2010, p. 10).

positioned over this sequence of locations than in the beneficiaries in W_1 . In an important sense, there are *more* beneficiaries in W_2 . (More on this below.)

	x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8	x_9	x_{10}	x_{11}	...
	p_1	p_2	p_3	p_4	p_5	p_6	p_7	p_8	p_9	p_{10}	p_{11}	...
W_1 :	1	0	1	0	1	0	1	0	1	0	1	...
W_2 :	1	1	1	1	0	1	1	1	1	0	1	...

From the assumptions above, there are (metaphysically) possible worlds W_3 and W_4 with the following local values, and which share every qualitative feature with W_1 and W_2 , respectively.

	x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8	x_9	x_{10}	x_{11}	...
	p_1	p_5	p_3	p_2	p_{15}	p_4	p_7	p_6	p_9	p_{10}	p_{11}	...
W_3 :	1	0	1	0	1	0	1	0	1	0	1	...
W_4 :	1	1	1	1	0	1	1	1	1	0	1	...

W_3 and W_4 differ from W_1 and W_2 by a slightly complex procedure: the persons at particular positions have been moved around; and some persons ($p_{15}, p_{35}, p_{55}, \dots$) obtain different values. Those persons do strictly better in W_3 than in W_2 , and strictly worse in W_4 than in W_1 . All other persons have the same local values in W_3 as in W_2 , and in W_4 as in W_1 . Assuming Pareto over persons, this means that $W_3 \succ W_2$ and $W_1 \succ W_4$.

But the persons in these worlds have been rearranged so that W_1 is qualitatively identical to W_3 , and W_2 to W_4 . So the pairs must have the same ranking: $W_1 \succcurlyeq W_2$ iff $W_3 \succcurlyeq W_4$.

Now suppose $W_2 \succcurlyeq W_1$. Then $W_2 \succcurlyeq W_1 \succ W_4 \succcurlyeq W_3 \succ W_2$. So this implies a violation of transitivity. So $W_2 \not\succeq W_1$.

But if we suppose that $W_1 \succcurlyeq W_2$, we can construct a pair of additional worlds W_5 and W_6 by a similar process such that $W_1 \succcurlyeq W_2 \succ W_3 \succcurlyeq W_4 \succ W_1$. That will lead to a violation of transitivity too. So $W_1 \not\succeq W_2$.

Therefore, in *Firing line*, W_1 and W_2 would be incomparable. Not equally good, but *incomparable*. But it was not so exotic a case. As it happens, we face a roughly analogous case whenever we influence infinitely many future persons, benefiting one set or another (with or without preserving their identities—see *ibid.*: 114–135). If we maintain Pareto over persons, incomparability would arise in (a close variant of) every case raised in the remainder of this paper.¹⁹

This is one of several reasons to reject Pareto over persons. With the assumptions above (which are hard to reject), it implies that we face incomparability in even mundane cases. It implies that the vast majority of the outcomes we produce are no better nor worse than the alternatives. And this is implausible. Our moral theory must say *something* in cases like *Firing line*.

¹⁹ The preceding argument is loosely adapted from Askill (*ibid.*), although Askill accepts this incomparability rather than abandon Pareto over persons.

On top of that, our moral theory must say the *right* thing which, at least according to intuition, is that it's worse to let every second person die than to let every fifth person die. (Or, for an analogous and even more compelling case, it would be worse to let every second person die than to let every *99th* person die!) Intuitively, the set of every second person seems larger than the set of every fifth: in particular, take any interval of three or more consecutive people and there are more of them. And of course the more people are killed, the worse the outcome is. And Pareto over persons rules that we cannot give this verdict. But Pareto over *positions* does not stop us, so I am quite happy to seek a \succsim relation which endorses Pareto over positions but not persons.²⁰

3.2 Impartiality

Here is another common objection to views like the one I'll present below. If we take the spatiotemporal position of value into account, we abandon a fundamental part of aggregative views—we fail to be *impartial*.

Impartiality is defended by aggregationists (Mill 1861; Sidgwick 1907; Singer 1972; Parfit 1986), in roughly the following form: it is morally no better for one to obtain the good than for another to. So too, it is no better for one group of persons—beneficiaries—to obtain the good than for another to. This applies whenever the amounts of the good obtained by the beneficiaries would be the same for either group, and the *number* of those beneficiaries is the same. And it applies regardless of *where* those beneficiaries are positioned and *who* the beneficiaries are.

But recall Finite Anonymity. If two worlds contain the same persons at the same positions, and they differ only by a rearrangement of any finite number of values from one set of beneficiaries to another, then Finite Anonymity (over *either* persons or positions) guarantees that the worlds are equally good. In practice, if we could either benefit *n* people by some amount(s) or *n* others by the same amount(s), and they start with the same amount of the good, then we should be indifferent between the two options.

Likewise, we might take two worlds with the same persons at the same positions and we rearrange their local values, and *also* rearrange which person is at which position. Still, we can rearrange any finite number of values and any finite number of persons and Finite Anonymity (over *either* persons or positions) guarantees that

²⁰ Despite the conflict between the two forms of Pareto, Pareto over persons does still have some intuitive draw. Given this, it may seem desirable to have $W_a \succsim W_b$ (or $W_a \succ W_b$) whenever Pareto over persons says so *and* Pareto over positions does not deny it. But Firing Line demonstrates that this can bring problems. In Firing Line, Pareto over persons says that $W_3 \succ W_2$ and $W_1 \succ W_4$ and Pareto over positions doesn't deny it, but that leads us to judging W_1 and W_2 as incomparable.

Given this, we must be cautious of following the verdicts of Pareto over persons even in narrow circumstances. But, I will briefly note, one situation in which it seems harmless to do so is when, across the two worlds we are comparing, only finitely many persons have their spatiotemporal positions differ between worlds. And my own proposal below satisfies Pareto over persons in those cases.

Another situation in which it seems harmless is when Pareto over positions (or some strengthening of it) says that two worlds are equally good. If so, and if Pareto over persons says that one world is strictly better, we might break the tie according to whatever Pareto over persons recommends. My own proposal doesn't do this, but it does seem a plausible modification to make.

the worlds are equally good. So it seems that, if \succsim satisfies Finite Anonymity (even just over positions), it preserves the impartiality we want.

Not all aggregation methods do satisfy Finite Anonymity. Discounting does not (e.g., Koopmans 1972). It involves summing local value in the usual manner, but only after *discounting* that value by some function of how far it is from the agent’s position. This will often allow us to compare infinite worlds, since a harsh enough discount rate can bring the total value down to a finite number. But it violates Finite Anonymity—it says that distant value is worth less (and very distant value almost *worthless*). It says that we should prefer to benefit our neighbors than to benefit distant strangers. Discounting may be a convenient solution to the problem, but I agree with Parfit (1986, pp. 480–486) and the others that it lacks the impartiality we want and so is morally indefensible. And that indefensibility is captured by its violation of Finite Anonymity.

But Finite Anonymity only requires indifference up to rearrangements of *finitely* many local values. What about rearrangements of *infinitely* many local values? Does impartiality require that we adopt a stronger principle of Finite Anonymity? No, it doesn’t. To see why not, consider these worlds. They’ll be familiar from above, but this time have the same persons (p_i) at the same positions (\mathbf{x}_i).

	p_a	p_b	p_c	p_d	p_e	p_f	p_g	p_h	p_i	...
	\mathbf{x}_1	\mathbf{x}_2	\mathbf{x}_3	\mathbf{x}_4	\mathbf{x}_5	\mathbf{x}_6	\mathbf{x}_7	\mathbf{x}_8	\mathbf{x}_9	...
W_1 :	1	0	1	0	1	0	1	0	1	...
W_2 :	1	0	0	0	1	0	0	0	1	...

W_1 is a rearrangement of W_2 —a rearrangement of infinitely many local values, over both persons and positions. We could take all of the 0-values in W_1 and bring them forward to occupy $\mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4$ and then $\mathbf{x}_6, \mathbf{x}_7, \mathbf{x}_8$, etc. And we could take the 1-values and spread them out to $\mathbf{x}_1, \mathbf{x}_5, \mathbf{x}_9, \dots$ and so on. If we were indifferent to even *infinite* rearrangements of value—if our evaluation was entirely independent of *who* and *where* the beneficiaries are—then we’d be indifferent between these two worlds.

But this is implausible. W_1 is clearly better, and Pareto (over either persons or positions) says so. So we cannot be wholly indifferent to infinite rearrangements of both persons and positions.

This doesn’t conflict with the core claim of impartiality: that it is no better to for one set of beneficiaries to obtain some values than another. It only holds when there is the same *number* of beneficiaries in either case. But there are infinitely many beneficiaries in both W_1 and W_2 . In response to this, we might say that the sizes of the sets of beneficiaries are simply undefined, so impartiality tells us nothing.

But we *can* define the sizes of infinite sets. How? We might do so by using their *cardinality*. Two sets have the same cardinality if and only if a bijection exists between them—for locations, if one set of locations can be rearranged (perhaps in infinitely many places) to precisely replace another set, just as we can do above with the 1-valued locations in W_1 and the 1-valued locations in W_2 . That case demonstrates that, in this *ethical* context, cardinality is not a plausible way to count beneficiaries. Impartiality would then force us to reject all forms of Pareto.

To satisfy Pareto while remaining impartial, we can turn to alternate notions of size for infinite sets. One is *containment*: a set A is larger than another, B , if $A \supset B$. That is, if A is a proper superset of B , and so *contains* it. And in the example above, the set of locations with value 1 in W_1 does contain the set with value 1 in W_2 . By containment, the set of beneficiaries is not the same size in both worlds—it's strictly larger in W_1 . And so impartiality will not imply, implausibly, that they are equally good. Nor must we violate Pareto.²¹

And there is another Pareto-compatible version of size as well—*density*. Compare, for example, the set of natural numbers to the set of even numbers. We might say that the natural numbers form the larger set because they contain the evens. But as we count along the naturals, from 0 up, it is also true that we encounter natural numbers more frequently than evens—in other words, the naturals have greater *density* over that ordered sequence. Likewise, in the example above, the locations with value 1 occur more densely over the sequence $(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \dots)$ in W_1 than they do in W_2 —in this sense, we can say there are *more* of them.

We can use density to measure subsets of locations more generally, when our locations have some essential properties which generate the right sort of structure. Whenever those properties are naturally represented in a coordinate structure (of any number of dimensions), we can say that one subset of locations occurs more densely over the space of possible values than another.²² For instance, the participants/positions in *Firing Line* above each had the property of where they were positioned in the line, and this was represented by a real number on a single coordinate. And we can say that the set of every second person occurred much more densely over that structure than the set of every fifth person. In general, if we treat spacetime positions as the relevant type of locations, they have an essential property which always generates such a structure: their coordinates in spacetime. So it can make sense to say that one infinite set of spacetime positions is larger than another if the first set occurs more densely over spacetime (see Wilkinson n.d.(a)). Alternatively, we might say that one set of *persons* is larger than another because it occurs more densely over spacetime, but this is a little strange—the size of the set could change if we moved those persons around, since a person's position isn't essential to their identity.

The key takeaway from all of this is that we have multiple candidate methods for measuring the size of a set of beneficiaries, not just cardinality.²³ And the others don't imply that we must be indifferent to rearrangements of infinitely many local values, since sets of locations need not be equally large just because we can biject them. To claim that impartiality *does* require indifference to such infinite rearrangements is to assume that cardinality is the relevant notion of size. And this requires independent justification, lest we beg the question.

So it is not at all obvious that views which are sensitive to the position of value in spacetime are immediately impartial, or that they constitute a fundamental departure

²¹ We might instead use a stronger notion of containment: a set A is larger than B if $A \setminus B$ has greater cardinality than $B \setminus A$.

²² This is an application of *asymptotic density* to this particular context (see Halberstam and Roth 1966).

²³ Another candidate is *numerosity*, as described by Benci and Di Nasso (2003), which is an extension of containment.

from aggregation. I see no compelling reason to disregard spatiotemporal views like mine, especially not when they give far more plausible verdicts (such as in *Firing Line* above).

4 Expansionism

If we give up Pareto over persons, and we don't ignore the spatiotemporal position of value, then we can compare many more pairs of worlds. Here is one way we might do so, by way of example.

Take worlds W_1 and W_2 (which resemble *Firing line*). Both worlds contain a sequence of persons spread over time, each holding the same position in both worlds. In W_1 , some harm befalls every second person. In W_2 , that harm befalls every fifth person. Those harmed bear value 0; those not harmed bear value 1. All other persons and positions in the world bear value 0.

We can depict these worlds on spacetime diagrams, as below. Each person is positioned in spacetime, so we can assign them some coordinates x in four dimensions. Here, the vertical axis represents time and the horizontal axis represents just the one spatial dimension, for simplicity (Fig. 1).

Which world is better? To determine this, let's gradually sum up the value in each world, but in a particular order. Start at some point P , as illustrated below. Sum the local values in the order of how far they are from P . For all locations within a given distance from P , those will lie within a circular region centered on P (or, in four dimensions, a hypersphere). To sum values in order of their distance from P is to sum them over some sequence of expanding circles of increasing radius (Fig. 2).

As we sum value over this sequence, we obtain a sequence of cumulative sums. These are the sums of local values within each circular region (the stage in the




Fig. 1 a W_1 , b W_2


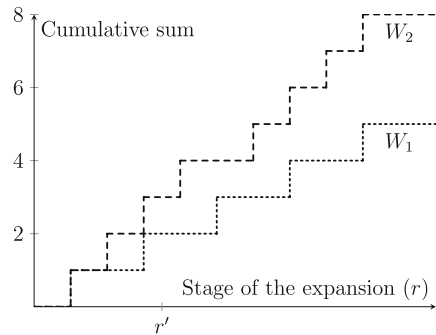


Fig. 2 **a** Expansions over W_1 , centered at P , **b** Expansions over W_2 , centered at P

Fig. 3 Cumulative sums in W_1 and W_2 , with expansions starting from P



expansion). We can represent these on a diagram too, as a function of the radius r of each circle (Fig. 3).

Notice that the cumulative sum for W_2 overtakes that of W_1 . There's some stage in the expansion (r') at which the cumulative sum of W_2 is higher than that of W_1 . And, beyond that stage, it remains higher.

Here, this happens no matter where we start the expansion. We could put the center of those circles in line with the values themselves (e.g., P), or somewhere off to the left or right (at Q or R in Figs. 3, 4). We could start at any real coordinates, and our expansion would still reach a stage at which W_2 has greater cumulative sum and, beyond which, W_1 won't ever catch up.

That's how we determine that W_2 is the better world: every sequence of expanding circles we could start from any point agrees that W_2 has more value. That's more value according to its cumulative sum, no matter how far out we take that sum.




Fig. 4 W_1

This is expansionism in a nutshell. More precisely, it’s a *spatiotemporal* version of expansionism—it uses the natural arrangement that locations (whether persons or positions) have in spacetime.

In general, expansionism consists in a principle like the following, from Vallentyne and Kagan (1997, p. 17).²⁴

SBI3 (Strengthened Basic Idea 3): W_1 is better than W_2 if

- (1) W_1 and W_2 have exactly the same locations, and
- (2) for all bounded regions of locations there is a bounded uniform expansion, such that, relative to all further bounded uniform expansions, W_1 is better than W_2 .

To clarify, a *bounded region* is best interpreted as a set of finitely many locations with non-zero value (*ibid.*: 13). An *expansion* is an infinite sequence of bounded regions, each a superset of the previous region. A *bounded uniform expansion* is an expansion, each region in which is obtained from the previous region by “...[adding] a band of constant width...” in all directions to the boundary (as in the example above) (*ibid.*: 16).

The principle I defend, and which was demonstrated in the example above, will deviate from Vallentyne and Kagan’s SBI3 in several important ways. Here is how, and why, it differs. (Readers uninterested in the motivation for these changes may skip ahead to the definition of SE1.)

²⁴ Arntzenius (2014) accepts a similar expansionist principle, designed to deal with cases of uncertainty. In cases of certainty, his principle is equivalent to SBI3.

(1) It uses the natural arrangement of locations in spacetime (whatever those locations are) to define regions and bounded uniform expansions, and so deals only with locations which *have* spatiotemporal positions.

Vallentyne and Kagan's principle does not necessarily use spatiotemporal positions to make judgments. It can be applied with other arrangements (and corresponding notions of boundedness and uniform expansions). For instance, using persons as locations, we could use the arrangement of those persons given by when we rank them by height. Or we might use their position in an alphabetical list of their names. In either case, we could start at any height (e.g., 170 cm) or any name (e.g., Xavier), and we could specify expanding 'regions' of persons over intervals of height or name (e.g., the interval [169 cm, 171 cm] or [Xander, Xena]). But we face problems in practice if we use such arrangements to make moral judgments. According to the physical theories cited above, our universe contains infinitely many tokens of every local phenomenon. This includes infinitely many persons who go by Xavier, Xander, or Xena, as well as infinitely many people with heights falling within any interval $[a - d, a + d]$. A non-zero proportion of them will have valuable lives, so these expansions would give an infinite (or undefined) cumulative sum at every stage. But not so if we use the spatiotemporal arrangement—for most plausible forms of moral value (such as pleasure), it is only physically possible to fit a finite amount of value into finite spacetime.²⁵ This is one reason to expand over the spatiotemporal arrangement. Another is simply that it seems far more plausible that moral evaluation depends on spatiotemporal position than that it depends on the heights, names, or other properties of persons or locations.

(2) The principle applies only if local value is additively separable.

This is an assumption which I expect Vallentyne and Kagan would agree with, but I want to make it explicit. It is also well-justified. For one, it is implied by all minimally aggregative views in the finite setting, and it is these are the views that we want to extend to the infinite setting. It is also very useful in practice—it allows us to simplify cases enormously, as we'll see in the following sections. Since local value is additively separable, we can add or subtract the same value to the same positions in each of two worlds, even infinitely many positions, and we wouldn't affect the ranking of those worlds. This means that, when comparing worlds, we can cancel out the local values that they have in common. For instance, if $W_1 \succcurlyeq W_2$, then $W_1 + W \succcurlyeq W_2 + W$ for all $W \in \mathcal{W}$, and vice versa.²⁶

(3) Expansions begin with points rather than regions.

This is equivalent to restricting initial regions, and expansions, to circles on the two-dimensional diagrams above—or, in four dimensions, hyperspheres—rather than any nasty old shape one can imagine. This is a costless strengthening of the principle—it allows us to judge in more cases without, as far as I can tell, giving

²⁵ See Wilkinson (n.d.(a)) for a full discussion of these points.

²⁶ Define addition of worlds as follows. The world $W_a = W_b + W_c$ is the world defined by $V_a(l) = V_b(l) + V_c(l)$ for all $l \in \mathcal{L}$ and $\mathcal{L}_a = \mathcal{L}_b = \mathcal{L}_c = \mathcal{L}$.

incorrect judgments in any. (Note that the problems described in the following section emerge with or without (4), but their solution requires it.) It is also in the spirit of Vallentyne and Kagan’s approach: their regions were already uniform in the way one expands to the next; now each region is uniform in its boundary as well.

In effect, this means that each region in the expansion is the set of all points within some fixed distance r of a central point P : a ball of radius r centered on P . We can abbreviate this as $E(r, P)$.

(4) The distance between locations is defined by geometric distance.

For expansions to be uniform, we need to define distance, which brings us to (5). The natural way to do this in the spatiotemporal setting is geometric distance d , given by $d^2 = \Delta x^2 + \Delta t^2$ for any two endpoints separated by spatial distance Δx and temporal distance Δt . On the two-dimensional diagrams used above, this corresponds to the distance given by placing a ruler on the page. And it results in each uniform expansion of a circle (or hypersphere) also being a circle (hypersphere).²⁷

(5) The principle provides a sufficient condition for equality (\simeq) between worlds, rather than just strict betterness.

Vallentyne and Kagan’s principle only stated a sufficient condition for strict betterness \succ . I want to add one for \simeq , along the same lines.

With those modifications in hand, here is what I take to be an initially plausible principle of spatiotemporal expansionism (SE).

SEI: For worlds W_1 and W_2 with the same set of locations,²⁸ $W_1 \succ W_2$ if, for all starting points P , there exists $r' \in \mathbb{R}$ such that for all $r > r'$,

²⁷ Geometric distance may seem the most natural way to do this, but it is not the only way. Since time and spatial distance measure fundamentally different quantities, and use different units, there are many different ways we might choose to combine the two into a common metric. But to allow us to construct useful expansions, such a metric does need to satisfy three conditions. (i) For every pair of points, the distance between them must be defined (or else expansions from one point will never reach the other point). (ii) If we start with a bounded region, any uniform expansion of that region must also be bounded (or else our cumulative sums may be infinite). And (iii) Distance must increase for increasing changes in space and increasing changes in time, so that the expansions become larger in the natural sense. Together, these conditions imply that our distance metric must satisfy $d^p = |\frac{\Delta x}{a}|^p + |\frac{\Delta t}{b}|^p$ for some positive, real a, b and n . (For instance, Arntzenius (2014: 44) describes an approach which effectively uses $a = b = 1$.) If we apply an additional condition—(iv) that regions of fixed distance are invariant under rotation—then we obtain geometric distance.

Note that, in a non-Euclidean space, there may be no definition of distance which satisfies (i)–(iii), let alone (i)–(iv), so we would need to abandon one of them. Wilkinson (n.d.(b)) sacrifices requirement (iii) to obtain a different distance metric. But even with the distance metric described there, we would obtain the same verdicts in the problem cases below. Likewise, if we adopted any metric which satisfies $d^p = |\frac{\Delta x}{a}|^p + |\frac{\Delta t}{b}|^p$, we would obtain the same verdicts, so we can simplify the discussion by using geometric distance.

²⁸ In practice, these sets will always be the same if we use positions as locations. If we use persons merely positioned in spacetime, they may not be.

$$\sum_{\mathbf{x} \in E(r, P)} V_1(\mathbf{x}) - V_2(\mathbf{x}) > 0$$

And $W_1 \simeq W_2$ if, for all P and all $r > r'$ the sum equals 0.

SE1 matches the method we used at the start of this section. When comparing any two worlds, start at any point you like (P). Generate larger and larger circles (hyperspheres)—circles with increasing r —centered at that point. Those are the regions $E(r, P)$. Take the cumulative sum of value within each of those circles (hyperspheres). Is there a stage at which one world takes the lead (where the difference in cumulative sums is greater than 0)? Does it remain in the lead for all greater r ? And does this happen no matter where you placed P ? If so, that world is better. Or is there a stage beyond which the cumulative sums remain equal? If so, the worlds are equally good.²⁹

So we have a way of comparing infinite worlds which isn't clearly implausible. In the remainder of this paper, I'll test SE1 on examples that demonstrate that it needs to be stronger. That stronger version will come in Sect. 6.

5 Problem cases

In the last section, I gave a seemingly plausible principle of spatiotemporal expansionism: SE1. Now, I'll test it on three fairly mundane cases, more complex versions of which would not be uncommon in our own universe. As it turns out, SE1 fails to give judgments in these cases; it falls silent.

Why is silence a problem? After all, SE1 is guaranteed to give us a \succsim relation which is incomplete. Otherwise it would violate Pareto or Finite Anonymity. But incompleteness comes in degrees—an incomplete relation may still be minimally complete, and deliver verdicts in cases which agents face in real life. And some agents will likely face cases resembling the three I describe below (absent some of my simplifications). So a principle which falls silent in these cases is unsatisfactory. This is why SE1 isn't quite enough, and why I'll strengthen it further in the next section.

5.1 Spatial shifts

Consider *Christmas*.

Example: Christmas

Each year the Consistent family celebrates Christmas. The Consistents always celebrate at their family home (at the exact same spatial location) and continue

²⁹ Note that SE1, and other expansionist approaches, can be expressed in terms of linearly ordered abelian groups and filters. Pivato (2014) shows this for all \succsim relations which satisfy basic conditions resembling Finite Anonymity and Pareto. I won't attempt to express SE1 in such terms here, but note that it is not necessarily inconsistent with approaches which use such representations (e.g., Bostrom 2011, pp. 20–24).

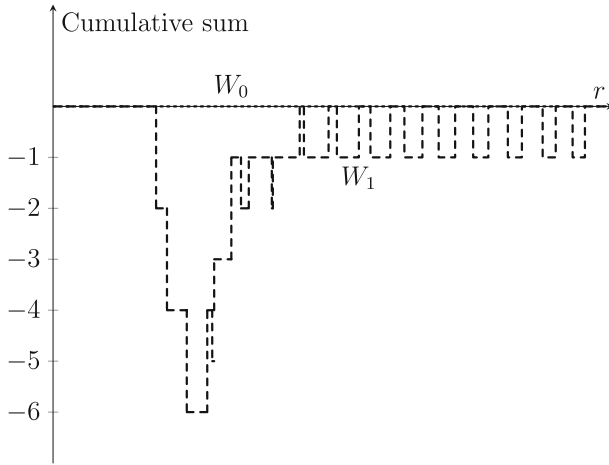


Fig. 5 Cumulative sums for W_1 and W_0 , with expansions starting from P

to celebrate at that home over successive generations. Celebrating Christmas brings joy to the Consistents, producing a brief, localised quantity of value. And the Consistents are lucky enough to have their family line persist forever, so there will be an infinite sequence of Christmases celebrated at their home. One generation of Consistents have the option of moving to a new address. If they move, every future Christmas will be celebrated at the new address. But they're a consistent family—they'll still celebrate at the same time, and produce the same amount of value. Everything else in the world will remain unchanged.

Staying at their current address can be represented by the zero world, W_0 , with local value 0 at all locations. The option of moving would then be represented by W_1 , pictured below. W_1 contains a sequence of local values -1 where they moved from, and values $+1$ where they moved to. And, since local values are additively separable, we ignore all other value in the world and just record the differences (Fig. 4).

Let's apply SE1. Starting our expansion at point P , we get the following cumulative sum for W_1 . W_0 has a cumulative sum of 0 at all stages of the expansion (Fig. 5).

Starting at P , our expansion always reaches the next -1 before its partner $+1$.³⁰ The expansion never reaches a stage at which the cumulative sum of W_1 remains greater than, less than, or equal to that of W_0 ; it will always alternate among these. And likewise if we start at point Q —the cumulative sum of W_1 will still alternate, but between greater than and equal to.

³⁰ This holds *even though* the delay between reaching the -1 point and the partner $+1$ point approaches 0. To see this, simply compare the distance from P to a -1 point to the distance from P to its partner $+1$ point—the latter is always strictly greater.

So SE1 is silent. If our betterness relation is fully described by SE1, these worlds remain incomparable. And that's a problem, because this case seems fairly mundane. We're just shifting a sequence of local value in space by a fixed amount. All the more mundane is that we didn't specify the distance we shifted it—it could be just a millimeter, and SE1 would still be silent. It seems plausible that agents really could make such an arbitrarily subtle change to the future, at least in addition to countless other changes, which would make the case all the more complicated. So I find it implausible that we cannot compare such worlds. Further, I find it implausible that these two worlds aren't *equally good*, as intuition suggests. Since SE1 cannot compare the two, nor justify our indifference between them, we have a problem.³¹ But worry not—we'll see below that it can be solved.

5.2 Temporal shifts

Consider *Public Holiday*.

Example: Public Holiday

A prophecy reveals that the kingdom of Alethkar will persist forever. To celebrate this, the queen will establish a new public holiday. This holiday will be accompanied by grand festivities and will be held on the same date, forevermore. Once it's established, its date will never be changed.

There are two dates to choose from, half a year apart. No matter when the holiday is held, it produces a fixed amount of value in every year except the first—in the first year, the delay will result in a slightly more enjoyable festival. The choice of date has no other effects.

The world in which the queen chooses the earlier date can be represented by the zero world W_0 . The world in which the queen chooses the later date is then W_1 (pictured below). W_1 differs from W_0 in two ways: there is local value -1 at every point the holiday would otherwise have happened (once per year); and there is local value 1 at every point the holiday happens instead, half a year later, with the exception of the first holiday in W_1 which is made slightly better by the delay. All of this happens at the same spatial position where the festivities traditionally take place (Figs. 6, 7).

³¹ Could SE1 be modified to avoid this verdict, while leaving intact its core features? Yes, but such modifications bring on other nasty implications.

One option is to abandon the requirement that expansions be *uniform*. For instance, we might replace them with sequences of square-shaped (hypercube-shaped) regions, as Askill (2018, pp. 208–11) does. This alternative may avoid the problem in *Christmas*, but it still faces the problems below. I should also note that this solution also appears somewhat arbitrary—it remains sensitive to *some* changes in Δx and Δt but not others.

Here is another option, indeed the only option if we hold fixed the uniform expansions and the dependence of \succsim on the cumulative sums over those expansions. When making judgments, do not require that *all* starting points agree on judgments. Instead of *supervaluating* over starting points, we could *subvalue*. After all, in the *Christmas* case, there is some starting point (R) for which the cumulative sums of W_0 and W_1 are equal for all expansions (and there are *none* which give a conflicting verdict). Why not then judge them as equally good? Upon examination, it turns out that doing so results in a \succsim relation which is intransitive.




Fig. 6 W_1




Fig. 7 Cumulative sums for W_1 and W_0 , with expansions starting from P

If we start the expansion at P , these are our cumulative sums:

This pattern repeats indefinitely. So, no matter how far we expand, there’s always a later stage at which W_1 has the lower cumulative sum, and one at which they’re equal. And this happens no matter where we start our expansion. So SE1 is silent in this case (as is SE1* from above)—these two worlds are incomparable.

This is odd. We only made a fairly mundane change to the world, one that agents might realistically achieve. But we still get incomparability? Again, we have a problem. (And again, as we’ll see, it can be solved.)

5.3 Chaotic effects

In practice, cumulative sums will almost always alternate, as in *Christmas* and *Public holiday*. To see why, consider *Writing or Netflix*.

Example: Writing or Netflix

Rita has a choice to make. She can either spend the next hour writing the rest of a paper on ethics, or spend that hour watching a television show. If she finishes




Fig. 8 W_1

the paper in the next hour and submits it for publication, it will be read by a policymaker. The policymaker will increase funding to some international aid intervention and thereby save k people from painful deaths. If Rita doesn't finish her paper within the hour, those k people will die. She knows all of this.

Unlike above, her actions have lasting, chaotic effects on the world. Saving those k people will, e.g., produce non-identity effects stretching indefinitely into the future (described in detail in Greaves 2016; Wilkinson n.d.(a)). If Rita takes one action rather than the other, this will cause random changes to many future events, continuing infinitely far into the future.

We can represent the actual outcome of Rita watching television by the zero-world by W_0 and the actual outcome of her finishing the paper by W_1 , as below. This world W_1 contains all of the differences between the two outcomes.³² These differences consist in: the initial benefit of saving k lives, proportional to k ; and an infinite sequence of random differences X_1, X_2, \dots (more on these shortly). For simplicity, I'm modelling the effects of Rita's actions as changing the value at one location per unit time, rather than at all times. I'm also pretending that these changes all occur at the same spatial position, rather than spread across her future as is more realistic. These simplifications won't change the verdicts below (Fig. 8).

Since Rita's actions make chaotic, seemingly random changes to the world, the differences X_i are given by *random variables* which are independent, identically distributed, and symmetric about 0. Those variables' probability distributions may reflect any of: Rita's subjective uncertainty about the exact effects of her actions; the evidential probabilities of the different outcomes; or their objective chances,

³² As above, this simplification is allowed by the assumption that value is additively separable.

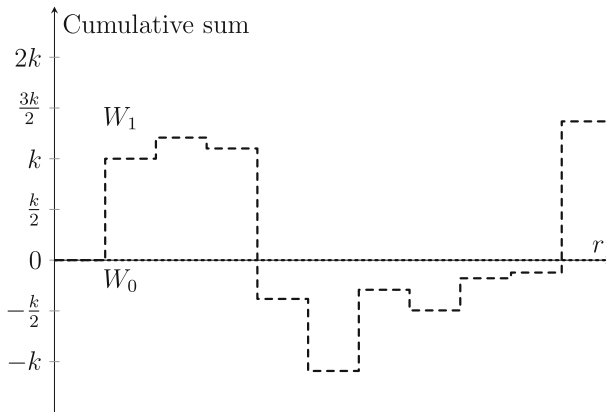


Fig. 9 Cumulative sums for W_1 and W_0 , with expansions starting from P

whichever is morally relevant. Either would work here. But note that, despite the mention of probabilities, we are not suddenly comparing *lotteries* over outcomes. This is still a comparison of *ex post* outcomes—of what would *actually* end up happening if Rita took a given action. But let’s pick out outcomes typical of this situation, given the uncertainty and chanciness present. We can do this by generating local values using the random variables X_i , and comparing whatever they give us. Due to the way they’re generated, these X_i will be random but still determinate. And, as we’ll see below, using these random variables allows us to draw much more general conclusions (Fig. 9).

To compare the outcomes, let’s run our expansions, starting from P (or any other point, which will give the same result). Here is our cumulative sum for a typical, randomly-generated W_1 .

Within this diagram, the cumulative sum of W_1 returns to 0 twice. But it will keep doing so: no matter how large r gets, it is guaranteed to return to 0 eventually. Why? The cumulative sum $(\sum_{\mathbf{x} \in E(r,P)} V_1(\mathbf{x}) - V_0(\mathbf{x}))$ forms a symmetric random walk over the reals, with starting point k and each step in the walk given by X_i . This random walk has an unfortunate property: it’s *recurrent* (Chungs and Fuchs 1951). For any r , no matter the value of the cumulative sum of W_1 at that stage, it will return to 0 again eventually. (It has probability 1 of doing so.) And this holds no matter how big k is.

What does this mean for *Writing or Netflix*? It means that SE1 won’t be satisfied for this W_1 and W_0 —it will remain silent. Nor will SE1 be satisfied in any situation like Rita’s with values generated the same way, with probability 1. And that means that, whenever our actions cause lasting and independent chaotic changes to the world, the outcomes of our actions are practically *guaranteed* to be incomparable by SE1.³³

³³ I am indebted to Christian Tarsney for raising this problem.

This is a problem. In real life, the actual effects of our actions will often be widespread, lasting and, to us, random. As an example (borrowed from Greaves 2016), suppose you are walking home from work and encounter an elderly person who needs help to cross a busy road. If you help them, both you and they will find yourselves at slightly different positions than you otherwise would have been at each time for the remainder of your journeys home. You will pass by many other pedestrians, momentarily advancing and delaying their journeys. When crossing the road itself, you momentarily delay several drivers. Given the dynamics of traffic, this momentarily delays countless other drivers (see Lighthill and Whitham 1955). And crucially, at least one of those numerous pedestrians and drivers would conceive a child that evening. By delaying or advancing their arrival home, you delay or advance the moment of conception and thereby change which sperm fertilises the egg. As a result, they conceive an entirely different child, who goes on to live a quite different life which, among other things, will change the moments of conception of countless others. These changes continue endlessly into the future and so are infinite in number.³⁴

Given this non-identity effect, my decisions will often resemble *Writing or Netflix*. Many of those local values in the future may be the same no matter what actions I take, but infinitely many others will be changed. So, many of the outcomes I must compare differ by a set of random local values. And SE1 won't be able to compare them.³⁵

6 Solution

Fortunately, all three problems are solvable. We can replace SE1 with a stronger, modified principle that gives all of SE1's judgments and more.

But first, to help make sense of that principle when I present it, consider a typical case in which SE1 says that $W_1 \succ W_2$ (illustrated below). For SE1 to say that, the graph of W_1 's cumulative sum must end up above that of W_2 at some point, and then stay above it for all greater values of r . The graph for W_2 might be above that of W_1 for some values of r (shaded with vertical lines), but only for finitely long—if we added together the lengths of every interval for which W_2 is in the lead, we'd get a finite length. Meanwhile, since W_1 will stay in the lead for all r above some finite

³⁴ To make matters worse, by helping the person cross the road and thereby altering the position of your body for a short while, you'll ever-so-slightly change the gravitational tidal forces acting on other astronomical objects. Over timescales of a more than 500 million years, this significantly alters the positions of planets in our solar system, and thereby changes the details of many events that happen thereafter (see Wilkinson n.d.(a)).

³⁵ This relates closely to the problem of simple cluelessness for objective betterness (see Greaves 2016). In the standard problem, one of two actions produces an initial benefit of k units of value, but we are clueless about the future impacts of both actions. Over long time horizons, the probability of either action turning out better approaches 0.5. So we are clueless of which is better. But in infinite cases like *Writing or Netflix*, by SE1, that probability drops to 0—we are practically guaranteed that the resulting worlds are incomparable.




Fig. 10 Typical cumulative sums for $W_1 \succ W_2$

bound, the total length that W_1 spends in the lead will be infinite. Note that the total area between the curves when W_1 is in the lead will be infinite too (Fig. 10).

Consider also some W_1 and W_2 which SE1 judges as equally good, as below. Beyond some finite value of r , their cumulative sums must be equal. So there can be at most a finite total length for the intervals during which W_1 is in the lead (shaded with horizontal lines) and likewise for the intervals during which W_2 is in the lead (shaded with vertical lines). Note also that the total area between the curves for the intervals during which W_1 is in the lead is finite, as is the total area for when W_2 is in the lead (Fig. 11).

We could generalize this to worlds which SE1 fails to compare. We could propose: take all the intervals of r values for which W_1 has the greater cumulative sum, and the intervals for which W_2 does; sum up the lengths of those intervals for each; if both total lengths are merely finite, then W_1 and W_2 are equally good; and, if the total length is infinite for W_1 and merely finite for W_2 , then W_1 is the better world.

This proposal would strengthen SE1 substantially. For instance, it gives us the correct judgment in *Christmas*. Below we have graphs of the cumulative sums for both worlds from that case. On those graphs, we frequently have intervals during which W_0 has greater cumulative sum (shaded with horizontal lines), but the lengths of the intervals become shorter and shorter as r increases. As it happens, the length of these periods shortens quickly enough that their total length converges to a finite sum. Meanwhile, there are no intervals during which W_0 has greater cumulative sum, so they have total length 0, which is finite too (Fig. 12).




Fig. 11 Representative cumulative sums for $W_1 \simeq W_2$

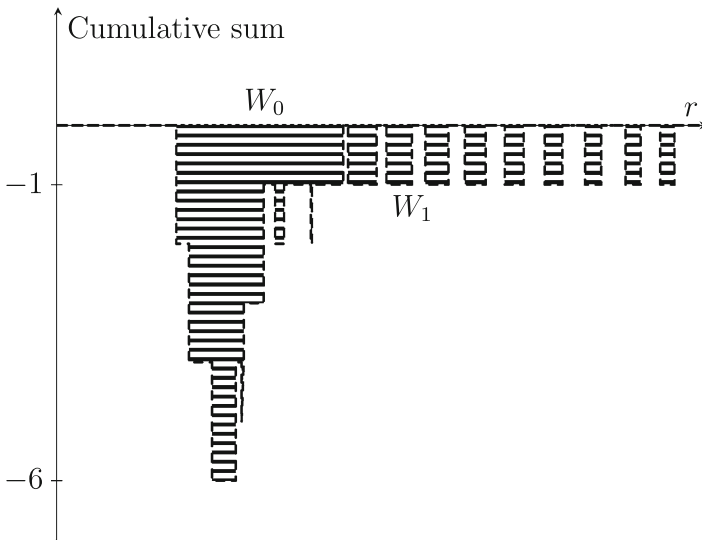


Fig. 12 Cumulative sums for W_1 and W_0 in *Christmas*, with expansions starting from P

And the same holds for all starting points in that case. For all of the intervals of r values during which either world has the greater cumulative sum, their total length is finite. So we might say that the two worlds in *Christmas* are equally good. And this is the correct answer, by intuition—the only difference between the two worlds is that we’ve shifted a sequence of values across in space slightly, which surely makes the world no better or worse.

But we’re still in trouble in *Public holiday*. As illustrated below, we have regular intervals of r values during which W_0 is in the lead (vertical lines) and regular intervals during which W_1 takes the lead. Since the cumulative sums repeat in this




Fig. 13 Cumulative sums for W_1 and W_0 in *Public holiday*, with expansions starting from P

pattern without end, both sets of intervals will have infinite total length. So we fall silent yet again (Fig. 13).

This seems odd. We may just as frequently have intervals during which W_1 leads, but its lead is always tiny. And recall from above that the very first holiday in W_1 obtained *slightly* higher value than the others. If we reduced its value to 1, like all the others, then W_1 's cumulative sum would never take the lead. We could then say that W_0 is better, but we cannot because of this tiny addition.

But we can overcome this. It seems important not only how frequently each cumulative sum is in the lead, but also how *great* that lead is. The occasional lead of W_1 (such as between r_2 and r_3) is one hundredth as big as the lead that W_0 takes (between r_1 and r_2), so it seems plausible that it should count for one hundredth as much. And this correlates exactly with the area between the curves during those intervals—the area of each interval during which W_1 is one hundredth the area of each interval during which W_0 leads. If we were to take a cumulative sum of the *area* between W_0 and W_1 , subtracting the area for intervals when W_1 leads, that sum would rapidly approach positive infinity. And likewise, for the exemplar comparison in Fig. 10, the area between the curves when W_1 leads minus the area when W_2 leads approaches positive infinity. Likewise for all cases in which SE1 says that one world is better than another. Meanwhile, in *Christmas* and the exemplar comparison in Fig. 11, that sum of areas would be merely finite.

We might thereby say which world is better in *Public holiday* and other cases, by considering the cumulative sum of *area* between the curves. In *Public holiday*, that sum rapidly approaches positive infinity (and it does so no matter where we put our starting point). So we could say that W_0 is better—the world in which the queen delays the public holiday is worse than the world in which she doesn't. This verdict may be counterintuitive, but I think it is the right one. After all, if the holiday were delayed by an entire year, then the world would only differ from the undelayed world in two ways: it would have one fewer holiday at the start (valued at -1), and a slight increase of value to the first holiday (0.01); the rest of the holidays would be

³⁶ SE2 can be stated more generally to allow for value spread densely over a region rather than concentrated at discrete points. To do so, we need only replace the sum with an integral: $\int_0^\infty \sum_{\mathbf{x} \in E(r,P)} V_1(\mathbf{x}) - V_2(\mathbf{x}) dr$.

identical between worlds. That would clearly be a worse outcome; roughly, it just involves removing value from the world. And since delaying a full year would be worse, it makes sense that delaying by less than a year would also make the world worse.

We can formalise this approach with SE2. Here, the sum on the right is the very same as what appeared in SE1: the difference in cumulative sums between W_1 and W_2 . But we want the area, so we need to multiply each of those sums by the length $r_{i+1} - r_i$ that it spends at that level. The start and finish of those intervals are given by $\{r_1, r_2, r_3, \dots\}$, the distances between the starting point P and each point \mathbf{x} at which there's a difference in value between W_1 and W_2 . And summing the areas between the curves over each such interval $[r_i, r_{i+1}]$, from r_1 out towards $+\infty$, we can see whether the total area is bounded or infinite.³⁶

SE2: Let W_1 and W_2 be worlds with the same set of locations. For any starting point P , let $\{r_1, r_2, r_3, \dots\}$ be the strictly increasing sequence of distances between P and each \mathbf{x} such that $V_1(\mathbf{x}) - V_2(\mathbf{x}) \neq 0$.

Then $W_1 \succ W_2$ if the following sum diverges unconditionally to $+\infty$.

$$\lim_{r \rightarrow \infty} \sum_{i=1}^r (r_{i+1} - r_i) \left(\sum_{\mathbf{x} \in E(r, P)} V_1(\mathbf{x}) - V_2(\mathbf{x}) \right)$$

And $W_1 \simeq W_2$ if the sum is bounded both above and below.

Based on the same reasoning as above, SE2 gives the verdict in *Public holiday* that it's better to schedule the holiday sooner rather than later. Likewise, it gives the same verdict in *Christmas*. And it implies SE1, so it implies all of the verdicts of it too, such as in *Firing line*. It also maintains the advantages of SE1 that it implies Finite Anonymity and Pareto over spacetime positions.

It also resolves our one outstanding problem case: *Writing or Netflix* (illustrated Fig. 14).

The problem here was that the cumulative sum of W_1 formed a symmetric random walk over the reals, which was guaranteed to be recurrent and so guaranteed to pass through 0 over and over without end. But the area between the cumulative sums forms no such random walk. Rather, it's guaranteed to *not* be recurrent—it's guaranteed to diverge unconditionally to $\pm\infty$. To see this, take any finite upper and lower bounds: it can be shown that the probability that the area under a symmetric random walk is within those bounds approaches 0 as $r \rightarrow \infty$ (Lipkin 2003). So the probability that the integral approaches $\pm\infty$ as $r \rightarrow \infty$ is arbitrarily close to 1. And the same holds no matter where we start the expansion. So W_1 is *guaranteed* to be comparable to W_0 by SE2. In fact, it's guaranteed that one of the worlds is strictly better.³⁷

³⁶ SE2 can be stated more generally to allow for value spread densely over a region rather than concentrated at discrete points. To do so, we need only replace the sum with an integral: $\int_0^\infty \sum_{\mathbf{x} \in E(r, P)} V_1(\mathbf{x}) - V_2(\mathbf{x}) dr$.

³⁷ Admittedly, there's a probability of $\frac{1}{2}$ that W_0 is the better world, so Rita would still be very uncertain of which act would turn out better. But this is just the same conclusion we face in the finite case (Greaves 2016).




Fig. 14 Cumulative sums for W_1 and W_0 in *Writing or Netflix*, with expansions starting from P

Thus, SE2 resolves all three cases. Spatiotemporal expansionism is back on firmer ground.

7 Conclusions

Infinite aggregation is hard to do. The standard approach of summing local values and comparing the totals won't suit our purposes in the infinite context. And that is disappointing for those of us who (i) hold minimally aggregative moral views, and (ii) most likely live in an infinite universe (as we do).

At first, spatiotemporal expansionism seems an odd view to adopt. It gives ethical significance to the spatiotemporal position of value. It even implies that it may sometimes be better not to make Pareto improvements to persons' wellbeing. I hope to have convinced you that the circumstances are dire enough, and the alternatives disturbing enough, that we should consider adopting it nonetheless.

In its most immediate form, spatiotemporal expansionism fails to give judgments in some very basic cases. This may seem to put it on a footing just as precarious as the alternatives, which entail widespread incomparability. But I suggest we adopt the tweaks I made in Sect. 6. Doing so, we can restore the judgments we wish to make in those cases. By doing so, we have a plausible way of restoring the aggregative judgments we wish to make in our universe, even if it is infinite.³⁸

Acknowledgements This paper has benefited enormously from the input of many brilliant and generous colleagues. I am indebted to Christian Tarsney, Peter Vallentyne, Teru Thomas, and especially Alan

³⁸ I've described how we can restore axiological judgments. I haven't proposed a way to restore our judgments of objective or subjective normativity. Nonetheless, extending this view to objective normativity is straightforward. And Wilkinson (n.d.(c)) provides a suitable extension to subjective normativity.

Hájek for each giving their frank and thorough feedback on drafts. I am grateful to Andreas Mogensen, Marc Fleurbaey, Adam Jonsson, Ralf Bader, Timothy L. Williamson, and to a seminar audience at the Princeton Philosophy Society for valuable discussion. On the logistical side, I am grateful the Princeton Department of Philosophy for allowing me to visit for much of the 2018–19 academic year, and the Australian-American Fulbright Commission for funding that visit, during which time this paper took shape. This work was also funded by both the Australian Government Research Training Program and the Forethought Foundation for Global Priorities Research, via a Global Priorities Fellowship; my thanks go to both funders for their support.

References

- Arntzenius, F. (2014). Utilitarianism, decision theory and eternity. *Philosophical Perspectives*, 28(1), 31–58.
- Asheim, G. B. d'Aspremont, C., & Banerjee, K. (2010). Generalized time-invariant overtaking. *Journal of Mathematical Economics* 46(4), 519–533.
- Askell, A. (2018). *Pareto principles in infinite ethics*. Ph.D. thesis, New York University.
- Bader, R. (2020). Person-affecting utilitarianism. In *Oxford handbook of population ethics*.
- Beckstead, N. (2013). *On the overwhelming importance of shaping the far future*. Ph.D. thesis, Rutgers University.
- Benci, V., & Di Nasso, M. (2003). Numerosities of labelled sets: a new way of counting. *Advances in Mathematics*, 173(1), 50–67.
- Bostrom, N. (2011). Infinite ethics. *Analysis and Metaphysics*, 10, 9.
- Cain, J. (1995). Infinite utility. *Australasian Journal of Philosophy*, 73, 401–404.
- Carroll, S. M. (2017). Why Boltzmann brains are bad. arXiv preprint [arXiv:1702.00850](https://arxiv.org/abs/1702.00850).
- Chung, K. L., & Fuchs, W. (1951). On the distribution of values of sums of random variables. *Memoirs of the American Mathematical Society*, 6, 12.
- Garriga, J., & Vilenkin, A. (2001). Many worlds in one. *Physical Review D*, 64(4), 043511.
- Gielis, J. (2003). *Inventing the Circle: The Geometry of Nature*. Antwerp: Geniaal Press.
- Gott III, J. R. (2008). Boltzmann brains—I'd rather see than be one. arXiv preprint [arXiv:0802.0233](https://arxiv.org/abs/0802.0233).
- Greaves, H. (2016). Cluelessness. *Proceedings of the Aristotelian Society*, 116(3), 311–39.
- Guth, A. H. (2007). Eternal inflation and its implications. *Journal of Physics A: Mathematical and Theoretical*, 40(25), 6811.
- Halberstam, H., & Roth, K. F. (2012). *Sequences*. Oxford: Clarendon Press.
- Hare, C. (2008). *The limits of kindness*. Oxford: Oxford University Press.
- Hare, C. (2010). Take the sugar. *Analysis*, 70(2), 237–47.
- Heyd, D. (2009). The intractability of the nonidentity problem. In M. Roberts & D. Wasserman (Eds.), *Harming future persons: Ethics, genetics and the nonidentity problem* (pp. 3–25). Dordrecht: Springer.
- Jonsson, A., & Voorneveld, M. (2018). The limit of discounted utilitarianism. *Theoretical Economics*, 13(1), 19–37.
- Knobe, J., Olum, K. D., & Vilenkin, A. (2006). Philosophical implications of inflationary cosmology. *British Journal for the Philosophy of Science*, 57(1), 47–67.
- Koopmans, T. (1972). Representation of preference orderings over time. In *Decision and Organization* (Vol. 57).
- Kripke, S. (1980). *Naming and necessity*. Oxford: Basil Blackwell.
- Lauwers, L., & Vallentyne, P. (2004). Infinite utilitarianism: More is always better. *Economics and Philosophy*, 20(2), 307–30.
- Lauwers, L. (2010). Ordering infinite utility streams comes at the cost of a non-ramsey set. *Journal of Mathematical Economics*, 46(1), 32–7.
- Lenman, J. (2000). Consequentialism and cluelessness. *Philosophy and Public Affairs*, 29(4), 342–70.
- Lewis, D. (1971). Counterparts of persons and their bodies. *Journal of Philosophy*, 68, 203–11.
- van Liedekerke, L. L., & Lauwers, L. (1997). Sacrificing the patrol: Utilitarianism, future generations and infinity. *Economics and Philosophy*, 13(2), 159–74.
- Lighthill, M. J., & Whitham G. B. (1955). On kinematic waves II. A theory of traffic flow on long crowded roads. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 229(1178), 317–345.

- Linde, A. (2007). Sinks in the landscape, Boltzmann brains and the cosmological constant problem. *Journal of Cosmology and Astroparticle Physics*. <https://doi.org/10.1088/1475-7516/2007/01/022>.
- Lipkin, L. (2003). Tossing a fair coin. *The College Mathematics Journal*, 34(2), 128–33.
- Mill, J. S. (1861 [1998]). *Utilitarianism*. New York: Oxford University Press.
- Parfit, D. (1986). *Reasons and persons*. Oxford: Oxford University Press.
- Pivato, M. (2014). Additive representation of separable preferences over infinite products. *Theory and Decision*, 77(1), 31–83.
- Schoenfield, M. (2014). Decision making in the face of parity. *Philosophical Perspectives*. <https://doi.org/10.1111/PHPE.1204>.
- Sidgwick, H. (1907). *The Methods of Ethics* (7th ed.). London: Macmillan.
- de Simone, A., et al. (2010). Boltzmann brains and the scale-factor cutoff measure of the multiverse. *Physical Review D*, 82, 6.
- Singer, P. (1972). Famine, affluence, and morality. *Philosophy and Public Affairs*, 1, 299–343.
- Svensson, L.-G. (1980). Equity among generations. *Econometrica*, 48(5), 1251–6.
- Temkin, L. S. (2012). *Rethinking the good: Moral ideals and the nature of practical reasoning*. Oxford: Oxford University Press.
- Vallentyne, P. (1993). Utilitarianism and infinite utility. *Australasian Journal of Philosophy*, 71(2), 212–7.
- Vallentyne, P., & Kagan, S. (1997). Infinite value and finitely additive value theory. *Journal of Philosophy*, 94(1), 5–26.
- Wald, R. M. (1983). Asymptotic behavior of homogeneous cosmological models in the presence of a positive cosmological constant. *Physical Review D*, 28(8), 2118.
- Wilkinson H. (n.d.(a)). Comparing infinite, chaotic outcomes. Unpublished manuscript.
- Wilkinson, H. (n.d.(b)). Infinite, invariant aggregation. Unpublished manuscript.
- Wilkinson, H. (n.d.(c)). Infinite aggregation and risk. Unpublished manuscript.
- Zame, W. (2007). Can intergenerational equity be operationalized? *Theoretical Economics*, 2(2), 187–202.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.