Every History

Jonathan Knutzen

Penultimate draft

Forthcoming in *The Philosophical Quarterly*

Abstract: This paper focuses on an underexplored challenge in infinite ethics. On realistic assumptions, if our universe is infinite, every nomologically possible history is actual and nothing we ever do makes a difference to the moral quality of the world as a whole. Call this thought *Every History*. This paper unpacks Every History and explores some of its ethical implications. Specifically, I argue that if Every History is true and the universe turns out to be infinite (1) our lives are globally insignificant, (2) moral principles enjoining the promotion of value need to be given an appropriately scope-restricted interpretation, and (3) impersonal consequentialism faces a serious challenge.

Key words: infinite ethics; ethical significance; meaning in life; consequentialism/nonconsequentialism; existential risk; longtermism.

1. Introduction

One relatively underexplored challenge in infinite ethics can be put briefly as follows.[1] If our universe is infinite, it contains all physically possible (finite) histories, and nothing we ever do affects which histories are actual. From these facts it appears to follow that, if our universe is infinite, nothing we ever do affects the moral quality of the world as a whole. Call this thought *Every History*. My aim in what follows is to spell out Every History in greater detail and explore some of its ethical implications.

[1] The challenge has not been much discussed, but see Knobe, Olum, and Vilenkin (2006) for a version of it.

There is a substantial literature in infinite ethics devoted to other challenges, like how to extend axiological and decision-theoretic principles from finite to infinite contexts.[2] Every History is, on its face anyway, not directly about these puzzles. It arises from a distinctive feature of physically realistic scenarios in an infinite universe: *parallel histories*. On plausible assumptions, an infinite universe contains a wide array of parallel histories. In some of these, events go one way; in others, they go differently. Given further assumptions (spelled out below) it seems we are powerless to change the overall pattern of value in the universe. This challenge can be appreciated without taking a stand on questions about value aggregation or decision-making under uncertainty in infinite contexts, matters about which I am agnostic. While there may be important connections between Every History and these other types of problems, I won't explore these here.

Infinity problems in ethics are rightly regarded as a significant threat to a certain consequentialist conception of ethics. Every History too seems threatening to that conception, though perhaps in a different way. As I'll explain below, it challenges views that take ethics to be chiefly about assessment and promotion of value from the radically de-centered perspective Henry Sidgwick (1907) memorably called the "point of view of the universe". In what follows, I refer to this type of view as *impersonal consequentialism*. My aim here is not to engage in a rescue operation aimed at preserving the spirit of impersonal consequentialism in infinite contexts. Instead, I treat Every History as a strong prima facie challenge to impersonal consequentialism and leave it open how best to proceed.

My own sympathies are nonconsequentialist. One of my goals in what follows is to illustrate the attractions of a conception of our ethical lives not centrally structured by the goal of value promotion when facing the prospect of an infinite universe. Yet the lesson of this paper is *not* that the prospect of an infinite universe makes consequentialism as such (i.e., a teleological moral theory that

[2] For a small sampling, see Arntzenius (2014), Askell (2018), Bostrom (2011), Meacham (2020), Vallentyne and Kagan (1997), and Wilkinson (2021).

defines rightness in terms of goodness) unattractive. Rather, it is that ethical theories will not be up to the task of helping us make sense of our ethical lives in an infinite universe unless they can accommodate a suitably *indexed perspective*—a perspective corresponding to our local history, a kind of "view from here". Consequentialist and nonconsequentialist moral theories alike have the resources to do this.

Two more brief clarifications are worth making at the outset. First, the structure of the argument is conditional. In the next section, I explain why Every History is plausibly true if the universe is (in relevant respects) infinite. I won't argue that the universe *is* (in relevant respects) infinite, other than to suggest that this appears to be a plausible interpretation of a range of cosmological models taken seriously by physicists. In subsequent sections, I argue that if Every History is true, this (plausibly) has some interesting ethical implications. The overall structure of argument, then, is doubly conditional: If the universe is infinite and if Every History is true, this has some interesting ethical implications. I won't argue that we have decisive reason to believe the antecedent of each conditional or pretend that the overall argument is rationally irresistible. But I do think we have reason to take the argument seriously.

Second, it is frequently considered a condition of adequacy for a view in infinite ethics that it can salvage at least modest claims about what we take ourselves to have reasons to do. I share this decidedly non-skeptical orientation. One might worry that Every History (and other puzzles) supports a more radical conclusion. Specifically, one might worry that Every History implies a kind of contingent nihilism about the significance of our actions according to which, if the universe is infinite, it doesn't matter what we do.[3] In my view, such nihilism is unwarranted for roughly Moorean reasons:

[3] Quentin Smith (2003) argues that if the future is infinite (and certain further assumptions hold), it doesn't matter what we do. He identifies this as a contingent form of nihilism. Eric Schwitzgebel and Jacob Barandes' (2024) recent argument that, in an infinite universe, almost everything we do causes almost everything (in some future galaxy) might be taken to support a similar conclusion. After all, if they are right, it pretty much doesn't matter what we do—take a breath, lift our pinky, fight for democracy—all our movements will give rise to countless goods and bads in the future and there's no way

it is much more likely that some step in the argument is mistaken than that, say, our reasons to pursue happiness or support people in need are illusory.[4] Relatedly, given its profound unattractiveness, nihilism should be viewed as an intellectual option of last resort. We should first see if we can make sense of our reasons within standard normative theories, making tweaks and modifications as needed, and accepting nihilism only if we must. As I hope to show in section 5, such ways of making sense of our reasons are available.

The paper is structured as follows. Section 2 begins by explaining Every History and motivating its plausibility. Section 3 argues that, in an infinite universe, our lives are globally insignificant: we make no difference, moral or otherwise, to what the world as a whole is like. Section 4 argues that, to avoid practical incoherence, moral principles enjoining the promotion of value need to be given an appropriately scope-restricted interpretation. And section 5 develops what is perhaps the most interesting upshot for moral theory, namely that Every History generates a serious challenge to impersonal consequentialism.

2. Every History

According to Every History, our universe contains all physically possible (finite) histories and nothing we ever do affects which histories are actual.[5] From this it follows that nothing we ever do affects the moral quality of the world as a whole. Every History thus consists of two interrelated hypotheses: one

to know which particular future effects our actions will have. Finally, and more directly analogous to Every History, Robert Adams (1979: 215-216) and David Lewis (1986: 123) entertain the worry that if modal realism is true, nothing we ever do matters.

[4] For example, Smith's (2003) argument is saddled with a series of strong and implausible assumptions. But even a more modestly framed argument for nihilism would, I think, face the difficulty that we have more reason to doubt one of its premises than to accept the conclusion that nothing we ever do matters.

[5] The restriction to *finite* histories is important. I do not explore whether our local history could be infinite toward the future or consider what ethical implications that would have. Assuming our universe (or pocket universe) is headed for inevitable heat death, it seems hard to believe that our actions could realistically have infinite future causal ripples. But see Schwitzgebel and Barandes (2024) for an argument to the contrary.

about the universe and one about us. Call the first of these the *every history hypothesis* and the second the *powerlessness hypothesis*.

Every history hypothesis: The universe contains all physically possible (finite) histories.

Powerlessness hypothesis: Nothing we ever do affects which (finite) histories are actual.

Let's explore these in turn.[6]

The universe's history consists of all the events that occur in it. This *global* history can be contrasted with the many *local* histories it contains. Our history—containing you and me and planet earth—is one such local history. If the universe is spatially infinite, it contains an infinite number of finitely sized local histories, histories in which other galaxies and planets form, other kinds of creatures evolve, and different events occur. For concreteness, we might imagine a local history as occupying a bounded region of space the size of our observable universe (or a "Hubble volume" of space) (cf. Garriga and Vilenkin 2001; Tegmark 2003). If the universe is spatially infinite, it contains infinitely many such local histories in distinct patches of space.

The every history hypothesis is a hypothesis of physical plentitude analogous to David Lewis's (1986) modal realism about possible worlds, except that the plentitude in question is realized in the actual world and constrained by the laws of physics (cf. Knobe, Olum, and Vilenkin 2006). The basic idea is that, if the universe is infinite, every nomologically possible finite history gets realized somewhere.[7] Some histories, e.g., those violating exact conservation laws, will never occur. But any history with a non-zero probability will occur. Since in an infinite universe the dice are cast an infinite

[6] As I'll point out at the end of this section, strictly speaking, the powerlessness hypothesis doesn't depend on the claim that, in an infinite universe, every nomologically possible history occurs. The powerlessness hypothesis could be derived from a weaker set of assumptions. Nevertheless, the every history hypothesis is plausible and serves handily for purposes of exposition.

[7] If we live in a multiverse with distinct local laws operating in distinct pocket or island "universes" (as in string theory's landscape multiverse), what is nomologically possible is, in an important sense, relative, viz., relative to a local regime of laws vs. relative to the most fundamental laws governing the entire multiverse ensemble. The set of possibilities within a local regime is obviously more constrained than the set of possibilities for the entire ensemble.

number of times at distinct locations, even the most unlikely scenarios are bound to occur. Given that there are infinitely many locations, every physically allowable finite history will occur somewhere.

The every history hypothesis appears to be an implication of a variety of cosmological models taken seriously in contemporary physics.[8] These include a low-density Friedman-Robertson-Walker universe (i.e., a model in which space extends forever without curving back on itself) (Ellis and Brundrit 1979); an inflationary model with infinitely many bubble "universes" (Garriga and Vilenkin 2001; Guth 2007; Knobe, Olum, and Vilenkin 2006; Vilenkin 1983); and the Everett interpretation of quantum mechanics (Carrol 2020; Wallace 2012). In the first, the histories play out in different contiguous regions of spacetime. In the second, they play out within distinct bubble universes (each of which is infinite) as well as across different bubble universes (i.e., at distinct locations within eternally inflating spacetime). In the last, they play out in different branches of the evolving wave-function in infinite-dimensional Hilbert space.[9]

Let's focus on the simplest of these for purposes of exposition: infinite contiguous space with low density and zero curvature. If space does not curve back on itself and has no "edge", it goes on forever. Assuming space is also filled uniformly (at large scales) with matter and energy, then there will be infinitely many finite local histories playing out across this infinite space (Ellis and Brundrit 1979; Knobe, Olum, and Vilenkin 2006; Tegmark 2003). The original configuration and subsequent evolution of particles in a given region of space is essentially stochastic. Within a given region of space, there is a randomly configured distributions of particles, which then evolve dynamically over time.

[8] I won't try to make the case for each model. For that (and nice overviews), see Ellis and Brundrit (1979), Knobe, Olum, and Vilenkin (2006), and Tegmark (2003).

[9] The every history hypothesis also appears to be an implication of the string multiverse and of cyclic universe models. Moreover, it may be an implication of any infinite universe in which Boltzmann brains/worlds of any arbitrary finite size, duration, and complexity are possible. For as long as there is a nonzero probability that a Boltzmann brain/world occurs, it will at some point and time eventually fluctuate into existence from surrounding low entropy states. Boltzmann brains/worlds appear to lead to skeptical conclusions and are (partly for that reason) controversial (Carroll 2017; Chen 2022).

Different original configurations (i.e., different "initial conditions") and how these dynamically evolve over time yield different histories at different locations. Since in an infinite universe the dice are cast an infinite number of times at distinct locations, even the most unlikely scenarios are bound to occur. Given that there are infinitely many locations, every physically allowable finite history will occur somewhere.

Precise repetition may or may not occur.[10] For our purposes, nothing rides on whether it does. It seems clear that, in an infinite space filled with matter and energy throughout, there will at least be *near* repetition of histories up to an arbitrarily high degree of precision. What matters for our purposes is near repetition that is sufficiently interesting in certain respects. In particular, we are interested in agents and their choices. In an infinite universe, there will be infinitely many copies of each agent participating in different local histories. Crucially, for our purposes, some histories will be identical or nearly identical up to a point and then diverge at a juncture of agential choice. Thus, an infinite universe will contain histories qualitatively identical or nearly identical to Earth's with copies of you and me deciding things differently and events subsequently unfolding differently from that point forward (Knobe, Olum, and Vilenkin 2006; Tegmark 2003). Call the set of qualitatively identical and similar individuals throughout the universe duplicates of each other.[11]

In our history, Gavrilo Princip assassinates Archduke Franz Ferdinand, which leads to World War I. In another history, Princip's duplicate does not assassinate the archduke's duplicate and World

[10] In classical physics, the states matter and energy can be in vary continuously so repetition appears not to be guaranteed. In quantum physics, according to at least one argument, there can only be finitely many "decoherent" histories in any finite volume of space so repetition of (coarse-grained decoherent) histories appears to be guaranteed (Garriga and Vilenkin 2001; Knobe, Olum, and Vilenkin 2006).

[11] Two brief comments about duplicates. First, in what follows, I sometimes refer explicitly to duplicates; at other times, I leave duplicate status implied, instead using a proper name or a personal pronoun to refer (indiscriminately) to both an individual and the individual's duplicates. Second, I take duplicates to be defined by similarity relations. As such, they are members of fuzzy sets. For a given fuzzy set of duplicates, it may be indeterminate whether some further individual is a member of that set. The resulting vagueness poses no deep problem. There are many individuals at other locations in the universe who will be sufficiently similar to you and me to count as our duplicates. It does not matter that there are many additional individuals such that it is indeterminate whether to count them as our duplicates.

War I never occurs. In our history, Princip chooses one way; in another, his duplicate chooses another way—with a very different outcome. Not all choices have such morally significant consequences, but some do. Given how human agents are constituted, they cannot realistically make any arbitrary choice. Not all options are genuinely "live" for agents, given their character and emotional dispositions (James 1897/1979). Sometimes agents are constrained, in Harry Frankfurt's terminology, by "volitional necessities" (Frankfurt 1988: 87). Thus, there might be choices Princip and his duplicates would *never* make. The decision about assassinating the archduke, though, was apparently not among them. We are interested here in those decisions that are genuinely "live" for agents.

The picture generalizes. For any choice an agent faces between morally relevant and live alternatives, X and Y, if the agent chooses X, some proportion of the agent's duplicates elsewhere in the universe choose Y. Put differently, whatever one does, the opposite will be done by some proportion of one's duplicates elsewhere in the universe.

To be sure, one can imagine a collection of histories all dramatically different and with nothing intermediate between them. Perhaps, for example, the qualitatively nearest alternative history to ours is one in which the Cambrian Explosion never occurs. There would then be no alternative histories in which Princip's duplicates choose differently. If the spectrum of histories is massively gappy in this way, then agents might not have duplicates who decide things differently elsewhere. But we have no reason to believe the spectrum of histories is gappy in this way. To suppose otherwise, one would need to suppose that some kind of selection mechanism operates on collections of particles to ensure that macroscopically described events are never similar, or never very similar in such a way as to diverge at a juncture of agential choice. We have no reason, I take it, to believe in such a cosmic selection mechanism. Histories are determined by the initial distribution and subsequent evolution of particles in a region of space. There may be a variety of macroscopically described "filters" (e.g., the emergence of self-replicating life, the emergence of an oxygen rich atmosphere and aerobic respiration,

etc.) that determine macroscopically similar trajectories. But we have no reason to rule out histories in which agents make different "live" choices.

Think of a moment of agential choice and its consequences in terms of a forking path: make one choice and the subsequent sequence of events goes one way; make a different choice and the subsequent sequence of events goes another way. In an infinite universe, each fork is taken by duplicates participating in different histories. Thus, for any morally relevant choice between live alternatives X and Y an agent faces, each of X and Y will be chosen somewhere in the universe by her duplicates. Since in an infinite universe every physically possible history is realized, so too is every physically possible history diverging at a juncture of agential choice.

We can now see why the powerlessness hypothesis is likely true in the kind of universe I've been describing. Your choice to X rather than Y does not change which histories occur. Nor does it shift the relative frequency of histories of a particular kind. Suppose you face the choice of helping Ayumi or not helping her. Suppose you decide to help her. If the universe is infinite, a duplicate of yours elsewhere in the universe will choose not to help Ayumi's duplicate. Do you make the world better by helping Ayumi? In one sense, yes. Had you chosen differently, Ayumi would now be worse off. In another sense, no. Whatever you choose, a duplicate of yours decides differently elsewhere. To capture this difference in your ability to be a difference-maker, we can call difference-making in our local history *local* difference-making, and difference-making to the moral quality of the world as a whole *global* difference-making. In an infinite universe, you can be a local but never a global difference-maker.

You could be a global difference-maker if you could affect the proportion of helping events to non-helping events throughout the universe (and, more generally, the overall pattern of histories). But you cannot do this. Suppose your duplicates help Ayumi's duplicates 2/3 of the time. What determines the ratio of helpings to non-helpings? Your local choice to help is *not* what explains the

ratio: that pattern is explained by the deep physical laws that determine relative frequencies of event types throughout the universe.

The moral quality of the world as a whole is determined by all the morally relevant features of the histories it contains. There are deep physical principles that determine the relevant patterning of histories. Your local choice does not change the global pattern. To make a difference to the moral quality of the world as a whole your local choice would need to be able to make a difference to the global pattern by affecting which histories are actual. But you cannot do that. Hence, you cannot make a difference to the moral quality of the world as a whole.

Call the totality of states of affairs relevant to assessing the goodness of a world its *evaluative profile*. For example, the amount of suffering a world contains is relevant to assessing how good or bad it is. Had our world contained either more or less suffering, it would have had a different evaluative profile. By the "moral quality of the world as a whole" I mean its total evaluative profile, the global pattern of value characteristic of our world.

Difference-making can be usefully understood in terms of counterfactual dependence.[12] Suppose that no matter how you act, the world's evaluative profile will be the same. Then you do not make a difference to the world's evaluative profile. To make a difference to the world's evaluative profile, you need to be able to affect which evaluative properties obtain. For example, suppose you help Ayumi and that if you don't nobody else will. This helping event is part of the world's total evaluative profile. Moreover, it seems things could have gone differently. Had you chosen differently, Ayumi would not have been helped. In this way, how you act locally appears to make a difference to the world's global evaluative profile. But while that is true in a finite world, if Every History is correct,

[12] The idea of counterfactual dependence has been influential in analyses of causation (Lewis 1973) and moral and criminal liability (Moore 2009). I make no claims either about causation or moral responsibility here. Note that, for all I say, there may be senses of causation in which the world's evaluative profile depends causally on your actions, and senses of moral responsibility in which you can bear some responsibility for it.

it turns out to be illusory in an infinite universe like ours. In such a world, your local choice does nothing to affect the global pattern of value.

The argument can now be summed up as follows. A wide range of cosmological models taken seriously in contemporary physics suggest that we live in an infinite universe. On realistic assumptions, this means the universe contains every nomologically possible finite history (*every history hypothesis*). Whether the world contains a single history or many histories, to make a difference to the moral quality of the world as a whole (i.e., its evaluative profile), your action must be able to make a difference to which history is, or which histories are, actual (i.e., the world's total evaluative profile must depend counterfactually on how you act). However, if the every history hypothesis is correct, then in an infinite universe you make no difference to which histories are actual. In such a universe, every nomologically possible history is actual and you cannot change which histories occur (*powerlessness hypothesis*). Hence, in an infinite universe, you cannot affect the moral quality of the world as a whole.

Notice that the powerlessness hypothesis doesn't strictly require maximal plentitude, i.e., the claim that absolutely every finite history that can happen will happen. Suppose only a large subset, say 95%, of all nomologically possible finite histories obtains. You still cannot change the global patterning of histories. Hence, the powerlessness hypothesis does not require the strong (and possibly controversial) assumption of maximal plentitude. Still, maximal plentitude is plausible and easy to grasp, so I've helped myself to it in the presentation of the argument and will continue to do so in what follows.

The scientific hypotheses referred to in this section are somewhat speculative, and the reasoning leading from them to the claim that we make no difference to the moral quality of the world as a whole might turn out to be flawed. Yet Every History cannot be lightly dismissed either. In what follows, I provisionally accept Every History and explore some of its ethical implications. I focus on three implications in particular.

3. Our lives are globally insignificant

The first and most general implication is that our lives are globally insignificant. Many of us want to make a moral difference, e.g., by giving to effective charities, fighting climate change, or campaigning for a good cause. Many of us also want to achieve things that make a nonmoral contribution, e.g., inventing something that improves people's lives or producing a great work of art. In engaging in these activities, we ordinarily take ourselves to be making a difference to what the world is like on the whole. However, Every History shows us that in an infinite universe this aspiration is misguided. Our lives cannot make any interesting difference, moral or nonmoral, to what the world is like on the whole. In an infinite universe, our lives are globally insignificant.

The idea of significance as difference-making is related to what Guy Kahane (2021) calls "importance". According to Kahane's Value Impact View,

> The degree to which something is important, relative to a domain, is a function of how much difference it makes to overall intrinsic value in this domain, compared to other things in it; and the more difference to value something makes, in this way, the more attention and concern it merits. (579)

According to Kahane, this view has implications for assessing the importance of human life within a cosmic context. "How important something is on the cosmic scale is a function of how much of a difference it makes to the overall intrinsic value of the entire universe, compared to other things" (Kahane 2021: 583). Hence, if there is no life elsewhere in the universe, or at least no life possessing the same valuable properties as humans, humans might enjoy great cosmic significance (Kahane 2014).

Kahane's view seems plausible, at least for some central usages of "significance" and "importance". However, this same view suggests that in an infinite universe, our lives are neither significant nor important. They make no difference to what the world is like on the whole (so they're globally insignificant). And when the entire cosmos is the relevant context, our lives merit no attention or concern vis-à-vis the rest of the universe, which contains infinitely many worlds populated by humans and similarly valuable forms of life (so they're globally unimportant).

Of course, even in an infinite universe, our lives will retain their local significance and importance, making a difference to how things go "around here". In that sense, they will be just as significant and important as we always thought they were. Moreover, mercifully, global insignificance and unimportance don't imply meaninglessness—not, at least, on plausible accounts of what meaningfulness amounts to. In particular, lives oriented around projects and relationships of genuine worth will continue to be rich in meaning (Wolf 2010). To be sure, insofar as significance contributes to meaningfulness (Kahane 2022), perhaps globally insignificant lives are less meaningful than globally significant ones. If that is right, then the available kinds and amounts of meaningfulness in an infinite universe may be reduced. Still, *most* of what makes life meaningful remains so even if our lives turn out to be globally insignificant.

4. Principles of value promotion need indexing

The second implication is that principles of value promotion need qualification, a point that is independent of the structure of moral theories incorporating such principles. Consequentialist views evaluate items (acts, institutions, rules, motives) as a function of value promotion; nonconsequentialist views reject this teleological structure. Nevertheless, nonconsequentialist views typically incorporate deontically constrained principles of value promotion. Specifically, many nonconsequentialist views contain imperfect duties of beneficence, which tell agents to do some amount of good while leaving

them discretion about how best to fulfill the requirement (Herman 2001; Kant 1785/1999; 1797/1999; Ross 1930/2002).

It is natural to read the injunction to promote value as unrestricted: to the extent agents are able, they ought to promote overall or global value in the world. However, Every History shows us that, if our universe is infinite, this is impossible. Facts about global value are fixed. Complying or failing to comply with the injunction to promote value in the world alters nothing about its global evaluative features. The aim of promoting global value, then, is not merely ambitious, specifying a goal we could hope to approximate; it is impossible.

To avoid enjoining an impossible aim, moral principles calling for the promotion of value need qualification. At least when reflectively endorsed in conjunction with the belief that the universe is infinite, the aim of world-improvement needs to be understood to be restricted in scope to our local history, i.e., to value promotion "around here". "Around here" need not be modest by conventional terrestrial standards. It could in principle include the entire "affectable universe", i.e., the portion of the universe we could, consistent with known physical laws, affect. From our ordinary perspective, this is a *huge* region of spacetime (Ord 2020: 232-233).

Every History nevertheless shows us that it is impossible to make an infinite universe better on the whole. To avoid practical incoherence, principles enjoining value promotion must be thought of as containing an indexical rider limiting value promotion to a cosmic patch of space centered on us (i.e., "Promote value around here!").

5. Impersonal consequentialism is inadequate

The final implication is that Every History poses a significant challenge to *impersonal consequentialism*. Recall that this is a view that combines teleological structure *and* commitment to assessing value from the radically de-centered perspective Henry Sidgwick (1907) called the "point of view of the universe".

Nothing rides on a particular phrase. Instead of speaking of the "point of view of the universe" we might talk of a "view from nowhere" (Nagel 1986) or, indeed, a "view from everywhere". The key idea—perhaps best captured by the last of these phrases—is that what matters for ethical assessment is all value bearing locations in the universe. It is this radically de-centered perspective that is the source of the challenge. As I noted in the introduction and will return to at the end of this section, the upshot is *not* that teleological moral views are in trouble; it is that reflecting on the challenge to ethics in infinite worlds reveals the need for a suitably indexed perspective, a "view from here".[13]

To illustrate the challenge to impersonal consequentialism, let's focus on maximizing welfarist act consequentialism. The right action, on this view, is the one that maximizes aggregate utility. The simplest way of interpreting this is in terms of total utility, i.e., the sum of all utility contributions at all times and places. In an infinite universe, of course, total utility remains the same regardless of what agents do. If our universe is infinite, then (realistically) it contains a countable infinity of goods and bads, including infinitely many good and bad lives.[14] Any action with finite consequences does nothing to change this. To preserve plausible aggregative verdicts (e.g., that saving five persons is better than saving one) in contexts where there are infinitely many utility contributions, some alternative to total utility is needed. There are a variety of proposals for how to do this.[15] Each has its attractions and downsides. Importantly, a set of impossibility results (Askell 2018; Lauers 2010; Zame 2007) show that various highly plausible principles cannot be jointly satisfied, so that *any* view about aggregation in infinite contexts is bound to yield some unsavory conclusions.

[13] Lewis (1986: 123-128) appears to draw a similar lesson, though in a different context, viz. a discussion of the implications of modal realism. Following J.J.C. Smart, Lewis thinks modal realism spells trouble for an understanding of consequentialism as concerned with value at every world. He instead favors an "ethics of our world" (128).
[14] There are conceivable universes with infinitely many utility contributions that contain only goods or only bads. However, any universe relevantly similar to ours will contain a mix of both goods and bads.
[15] See references in footnote 1.

As I noted in the introduction, I assume no particular view about aggregation in infinite contexts here. What I do assume, for the sake of argument, is that one or other proposal works. Given this assumption, I assume there is a way to formulate aggregative axiology that preserves the spirit of impersonal consequentialism in infinite contexts. There is still, I think, an unaddressed problem about whether we can really make sense of our ethical lives in the way impersonal consequentialism requires if our universe is infinite.

To illustrate the problem, consider a simple case encountered previously. Ayumi needs help and you are in a position to help her. Let's add that you can help at very low personal cost and that there are no competing moral demands on you. Should you help? The obvious answer is yes. Let's assume, then, that you have most reason to help. The question is which type of ethical theory best makes sense of that verdict. One possibility is to view your reasons in nonteleological terms. Perhaps it is simply appropriate that you respond to Ayumi by helping her: it is an apt response to her as a being of a certain sort, not a way of promoting valuable states of affairs. That interpretation seems quite plausible to me, and I'll have more to say about this way of understanding your reasons shortly. Another possibility is to view your reasons in teleological terms: you have reason to help Ayumi because doing so would bring about a valuable state of affairs. The unaddressed problem arises from combining this teleological conception of your reasons with a "view from nowhere" conception of ethics.

The question is why it matters how things go in the particular history you are located in. Given Every History, you know that every possible history will occur. In some histories, Ayumi will be helped; in others, she will not be helped. You can change nothing about which histories are actual and make no difference to the moral quality of the world as a whole. So why does it matter whether you happen to be located in one of the histories where Ayumi is helped? More generally, reflective moral agents who come to believe that the universe is infinite will realize that no choice of theirs ever makes

a difference to the overall value or disvalue of the world. Given Every History, they will realize that their choices are, in fact, self-locating rather than world-determining. Rather than making a difference to the moral quality of the world as a whole, their choices merely locate them within a collection of qualitatively similar histories.[16]

The issue is not whether you have most reason to help Ayumi: we assume that you do. Hence, you ought to help Ayumi. The issue is also not about whether a suitable aggregative principle for infinite contexts will deliver the verdict that helping Ayumi is best: we assume that there is such a principle. Hence, it is possible to spell out in teleological terms why you have most reason to help Ayumi. Finally, the issue is not about what would be most virtuous: we assume that helping Ayumi is the virtuous thing to do. Hence, if you are appropriately morally motivated, you have reason to want to be located in a history where you help Ayumi.

The issue, instead, is about which ethical commitments make sense of your situation as an agent faced with the prospect of doing good locally. To make sense of the thought that it matters how you act in a local history, you need some sort of indexed perspective. But it is hard to see how impersonal consequentialism can provide such a perspective. Suppose, for concreteness, that you should help Ayumi because Pareto (over persons or locations) is true.[17] Then Pareto (over persons or locations) explains why you have most reason to help Ayumi assuming a teleological conception of your reasons. But by hypothesis, copies of yours decide things differently elsewhere in the universe, and all possible histories play out somewhere. Impersonal consequentialism attaches no special weight to you or to your location. All that matters is value from the point of view of the universe. From that

[16] Since the "same" choice is made in qualitatively distinct histories, a given choice does not locate an agent within any qualitatively unique history but instead within a collection of qualitatively similar histories (viz., those in which the choice is made).

[17] Pareto over persons says, roughly, that you should help Ayumi because you make one agent better off and no one else worse off; pareto over locations says, roughly, that you should help Ayumi because doing so makes things go better in one region of the universe without making things go worse in any other. For Pareto over persons, see Askell (2018). For Pareto over spacetime locations, see Wilkinson (2021).

de-centered perspective, it does not matter where you happen to be located. From that perspective, your local choice to help or not changes nothing globally. Suppose you decide *not* to help. Then you make nothing worse globally. And, given Every History, you know this. You know that if you don't help, you are located in one of a collection of similar histories in which you and your copies do not help Ayumi. But you also know that there are other histories in which your copies do help Ayumi. Indeed, you know that all possible histories occur, with some ratio of helping to non-helping events and, more broadly, with various aggregate sums of value within distinct (finite) local histories. The only question is which of these histories you are located in. From a de-centered perspective concerned with aggregate value in the universe, it doesn't matter which one of these histories you happen to be located in. How could it? It might matter *to you* because you are virtuous and care about doing what you have most reasons to do. But from a point of view concerned solely with aggregate value in the universe, clearly it doesn't matter which history you happen to be in.

To further illustrate the issue, consider the ethics of existential risk. Many of us find it important that humanity survives extinction in the coming centuries and beyond and believe ourselves to have reasons to ensure that it does so (Bostrom 2013; Ord 2020; Scheffler 2013; Schubert et al. 2019). But *why* is humanity's survival important? Ethical longtermists have recently argued that we should think about this question in terms of the expected value of humanity's long-term survival (Beckstead 2013; Greaves and McAskill 2019; McAskill 2022). Longtermists believe that what we ought to do is mostly determined by the long-term consequences of our actions. The simplest way of supporting longtermism and arguing for the importance of humanity's long-term survival is in terms of the idea that we ought to maximize aggregate utility.

In an infinite universe, however, there will be infinitely many histories containing duplicates of earth and its denizens. Some of these go well by way of maximizing utility in that history (i.e., locally); others go badly. In some histories, humanity survives for billions of years, seeding the cosmos

with life and giving rise to advanced intelligent life forms. In others, humanity goes extinct in the late Pleistocene. And in many others, humanity goes extinct at some intermediate point. All these histories play out in different regions of spacetime. From an impersonal perspective scoped out to the entire universe, it is not clear why we should care whether the history we happen to be part of is one in which humanity goes extinct, say, early in the 22rd century or late in the 567th. The "point of view of the universe" attaches no special importance to our location. Since it attaches no special importance to our location, there seems to be no special reason to care about how things go here, i.e., in our particular history.

The simplest axiology underwriting longtermism is risk neutral totalism, a view that is additive, impartial, and risk neutral. Christian Tarsney and Hayden Wilkinson (2022) argue that a variety of proposals about how to extend risk neutral totalism to infinite contexts are plausible, and that all these extensions in essence allow agents to ignore the portion of the universe they cannot affect and focus exclusively on the portion they can affect. If they are right, there exist plausible aggregative principles on which it is perfectly coherent to claim that acting to ensure humanity's survival in an infinite universe maximizes expected value (without changing the sum total of value in the universe). Suppose that's right. There seems to be a further question, not settled by the existence of a coherent aggregative principle for expressing the claim that, even in an infinite universe, working to ensure humanity's survival maximizes value. Why should we care how long humanity survives in the particular history in which we happen to find ourselves? In an infinite universe, humanity will have many different lifespans, yielding distinct utility contributions in different regions of spacetime. From a radically de-centered perspective, it doesn't seem to matter what humanity's lifespan happens to be (i.e., how much utility there is) in the particular history we are located in.

Let's briefly explore the resources of a very different ethical conception to provide a suitably indexed perspective. According to nonteleological conceptions of our reasons, things of value demand

to be engaged with and responded to in particular ways, ways that cannot readily be cashed out in terms of what promotes value (Anderson 1993; Scanlon 1998). While such a view is not limited to *moral* reasons, it can be grasped with special vividness in the moral domain. For example, humans are owed respectful treatment. They must not be enslaved, tortured, or killed, and they ought to be provided with sufficient resources to live minimally decent lives. To fail to treat them in these ways violates their rights and fails to accord them the kind of respect they are due. The special value of humans grounding these and other rights is often called their dignity (Griffin 2008; Nussbaum 2011). On a nonteleological conception, reasons to adequately respect the dignity of persons are not merely reasons to promote valuable state of affairs; they are reasons to respond appropriately to the value persons have. Facts about the size of the universe or about the location of agents in it seem irrelevant to such nonteleological reasons. Persons are owed respectful treatment no matter what happens in distant regions of space or where we happen to be located. In an infinite universe, slavery and torture are just as inappropriate as they are in a finite universe; they ought never occur. This is true, even if, regrettably, in an infinite universe, there are infinitely many instances of slavery and torture. A nonteleological conception of reasons, then, arguably provides an indexed perspective keyed to local engagement with value.

Personal reasons provide an additional resource. On a nonteleological conception of reasons, personal reasons are a special case of nonteleological reasons more generally. However, it is worth discussing them separately, since personal reasons provide an additional resource for articulating an indexed perspective. Frequently discussed examples of personal reasons include reasons arising from personal projects and loving attachments (Scheffler 1982; Williams 1981). Because they are agent-relative, personal reasons add a further layer of perspectival indexing to the kinds of agent-neutral reasons considered above. Take two examples from Joshua Knobe, Ken Olum, and Alexander Vilenkin (2006). The first involves a father who feels concern for his daughter. The father runs to

comfort his daughter when she cries. Nothing changes when he comes to believe that there are other sequences of events elsewhere in the universe involving duplicates of him and his daughter where things go differently. The second involves self-interest. Suppose you come to believe that there are duplicates of yours in other regions of the universe who have either better or worse health outcomes than you. Should this affect how you act or feel? Obviously not. Prudential concern is not threatened in the slightest by finding out that the universe is infinite. What these examples show is that the sphere of personal concern—whether focused on self or others—provides powerful reasons to want to be located in a history in which things go well for the object of one's concern, i.e., to be a local difference-maker.[18]

To round things out, consider what a nonteleological conception of our reasons to ensure human survival might look like. There are a variety of such conceptions. However, consider two recent proposals that exemplify the kinds of nonteleological reasons just described. Johann Frick (2017) argues that we have reasons to ensure humanity's persistence because humanity has final value. On Frick's understanding, the relevant reasons are nonteleological. They do not derive from the fact that preserving humanity would make the world any better. Instead, they derive from the fact that preserving humanity is, all else equal, the fitting response to humanity's final value. Samuel Scheffler (2018) articulates a complimentary but slightly different perspective. According to Scheffler, among our reasons for wanting humanity to survive under conditions hospitable to flourishing are reasons of interest and reasons of love (2018: 40-67). Reasons of love are similar to those Frick discusses under the heading of reasons to respond appropriately to final value; they are reasons to care about human survival for its own sake (under conditions hospitable to flourishing). Reasons of interest, by contrast, add an interesting relational twist to the story, one that makes ineliminable reference to our own lives, projects, and commitments, and to the particular cultures, institutions, traditions, and practices our

[18] See Lewis (1986: 125-126) for a similar point with respect to the implications of modal realism.

lives are embedded in. From this relational perspective, we appropriately care for humanity and its future in a way that is analogous to our concern for ourselves, our children, our romantic partners, our friends, and our associates. The relevant reasons do not derive from the fact that preserving humanity would make the world any better but from the fact that preserving humanity is an appropriate response to a form of final value with which our own lives and projects are deeply bound up. Neither Frick's nor Scheffler's account of our reasons to preserve humanity seems especially vulnerable to worries about infinite parallel histories.

Nonteleological conceptions of ethics, then, appear to be well-equipped to supply an indexed perspective. These conceptions, in effect, tell agents to engage appropriately with value. While there's obviously an important substantive question about what that entails from one case to the next, it always means: engage appropriately with *local* value! What else could it mean? The perspective is automatically keyed to engagements of finite agents with local occurrences of value.

In sum, impersonal consequentialism eschews all local perspectives and gives agents no reason to care about how things go in the particular histories they are part of. By contrast, nonteleological conceptions of ethics provide a natural indexed perspective keyed to local engagement with value. Let me make two concluding observations about the upshot of these observations. First, as I've said, the upshot is that ethics in an infinite universe needs a suitable indexed perspective, not that views with a teleological structure are unattractive. The prospects for teleological conceptions to accommodate an indexed perspective seem good. There are already a variety of proposals for incorporating an agent-centered perspective within consequentialist moral theory (e.g., Howard 2022; Portmore 2007). Plausibly, an analogous perspectival accommodation could be made to provide the needed indexed perspective for consequentialist moral theories in an infinite universe. To fans of aggregative ethics, my message is not that aggregative principles for infinite contexts are in trouble. It is that whatever the correct set of aggregative principles turns out to be for such contexts, they will have to be

accompanied with the clear-eyed recognition that, if our universe is infinite, we cannot be global difference-makers and all we can do is maximize value in the particular local history we are part of. Second, I don't take the challenge to impersonal consequentialism to necessarily be decisive. I think it is a strong challenge, and it should make us doubt whether impersonal consequentialism is up to the task of guiding our ethical lives in an infinite universe. But, of course, all ethical theories face challenges. Whether this particular challenge is decisive depends on an assessment of the theoretical costs and benefits of all relevant contenders. That's not something I've undertaken (much less accomplished) here. Hence, for all I've argued, it is possible that impersonal consequentialism is the best supported moral theory, even in an infinite universe.

6. Conclusion

The possibility that the universe is infinite raises important puzzles for ethics. My focus here has been on a problem that has received relatively little attention. According to Every History, in an infinite universe, every nomologically possible history is actual and nothing we ever do makes a difference to the moral quality of the world as a whole. I argued that, on realistic assumptions about what it would mean for our universe to be infinite, Every History is plausibly true. I then explored three implications. The first is that, in an infinite universe, our lives are globally insignificant. The second is that, to avoid practical incoherence, moral principles enjoining value promotion must be given an appropriate scope-restricted interpretation. The third is that ethics needs some kind of indexed perspective corresponding to agents' location within particular histories, something impersonal consequentialism seems unable to provide.[19]

[19] I am grateful to two anonymous referees for this journal for their excellent comments. I would also like to thank Eric Schwitzgebel and Travis Timmerman for helpful feedback on an earlier draft, and Cian Dorr and Eddy Keming Chen for helpful discussion.

References

Adams, R. (1979) 'Theories of Actuality', in Michael J. Loux (ed.), *The Possible and the Actual: Readings in the Metaphysics of Modality*, 190-209. Ithaca: Cornell University Press.

Anderson, E. (1993) *Value in Ethics and Economics*. Cambridge, MA: Harvard University Press.

Arntzenius, F. (2014) 'Utilitarianism, Decision Theory and Eternity', *Philosophical Perspectives* 28 (1), 31–58.

Askell, A. (2018) *Pareto Principles in Infinite Ethics*. PhD thesis, New York University.

Beckstead, N. (2013) *On the Overwhelming Importance of Shaping the Far Future*. PhD thesis, Department of Philosophy, Rutgers University.

Bostrom, N. (2011) 'Infinite Ethics', *Analysis and Metaphysics* (10), 9-59.

Bostrom, N. (2013) 'Existential Risk Prevention as Global Priority', *Global Policy*, *4*(1), 15-31.

Carroll, S. (2020) 'Why Boltzmann Brains are Bad', in S. Dasgupta, B. Weslake & R. Dotan (eds.), *Current Controversies in Philosophy of Science*. Routledge.

Carroll, S. (2020) *Something Deeply Hidden: Quantum Worlds and the Emergence of Spacetime*. Penguin.

Chen, E. K. (2022) 'Time's Arrow and Self-locating Probability', *Philosophy and Phenomenological Research*, *105*(3), 533-563.

Ellis, G., and G. Brundrit (1979) 'Life in the Infinite Universe', *Quarterly Journal of the Royal Astronomical Society*, *20*, 37-41.

Frankfurt, H. (1988) 'The Importance of What We care About', in *The Importance of What We Care About: Philosophical Essays*. Cambridge: Cambridge University Press.

Frick, J. (2017) 'On the Survival of Humanity', *Canadian Journal of Philosophy*, *47*(2-3), 344-367.

Garriga, J., and A. Vilenkin (2001) 'Many Worlds in One', *Physical Review D*, *64*(4), 043511.

Greaves, H., and W. MacAskill (2019) 'The Case for Strong Longtermism', *GPI Working Paper No. 7-2019*.

Griffin, J. (2008) *On Human Rights*. Oxford: Oxford University Press.

Guth, A. (2007) 'Eternal inflation and its implications', *Journal of Physics A: Mathematical and Theoretical*, *40*(25), 6811.

Herman, B. (2001) 'The Scope of Moral Requirement', *Philosophy & Public Affairs*, *30*(3), 227-256.

Howard, N. (2022) 'Consequentialism and the Agent's Point of View', *Ethics*, *132*(4), 787-816.

James, W. (1897/1979) *The Will to Believe and Other Essays in Popular Philosophy*. Cambridge, MA: Harvard University Press.

Kahane, G. (2014) 'Our cosmic insignificance', *Noûs*, *48*(4), 745-772.

Kahane, G. (2021) 'Importance, Value, and Causal Impact', *Journal of Moral Philosophy*, *19*(6), 577-601.

Kahane, G. (2022) 'Meaningfulness and Importance', in I. Landau (ed.), *The Oxford Handbook of Meaning in Life*. Oxford: Oxford University Press.

Kant, I. (1785/1999) '*Groundwork of the Metaphysics of Morals*', in M. Gregor (ed.), *The Cambridge Edition of the Works of Immanuel Kant: Practical Philosophy*. Cambridge: Cambridge University Press.

Kant, I. (1797/1999) '*The Metaphysics of Morals*', in M. Gregor (ed.), *The Cambridge Edition of the Works of Immanuel Kant: Practical philosophy*. Cambridge: Cambridge University Press.

Knobe, J., K. Olum, and A. Vilenkin (2006) 'Philosophical Implications of Inflationary Cosmology', *The British Journal for the Philosophy of Science*. (57)1, 47-67.

Lauwers, L. (2010) 'Ordering Infinite Utility Streams Comes at the Cost of a Non-Ramsey Set', *Journal of Mathematical Economics*, 46(1), 32–7.

Lewis, D. (1973) 'Causation', *The Journal of Philosophy*, *70*(17), 556-567.

Lewis, D. (1986) *On the Plurality of Worlds*. Blackwell Publishing.

MacAskill, W. (2022) *What We Owe the Future*. Basic Books.

Meacham, C. (2020) 'Too much of a good thing: Decision-making in Cases with Infinitely Many Utility Contributions', *Synthese* 198, 7309-7349.

Moore, M. (2009) *Causation and Responsibility: An Essay in Law, Morals, and Metaphysics*. Oxford: Oxford University Press.

Nagel, T. (1986) *The View from Nowhere*. Oxford: Oxford University Press.

Nussbaum, M. (2011) *Creating Capabilities: The Human Development Approach*. Cambridge, MA: Harvard University Press.

Ord, T. (2020) *The Precipice: Existential Risk and the Future of Humanity*. London: Bloomsbury Publishing.

Portmore, D. (2007) 'Consequentializing Moral Theories', *Pacific Philosophical Quarterly*, *88*(1), 39-73.

Ross, D. (1930/2002) *The Right and the Good*. Oxford: Oxford University Press.

Scanlon, T. (1998) *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.

Scheffler, S. (1982) *The Rejection of Consequentialism*. Oxford: Oxford University Press.

Scheffler, S. (2013) *Death and the Afterlife*, edited by N. Kolodny. Oxford: Oxford University Press.

Scheffler, S. (2018) *Why Worry About Future Generations?* Oxford: Oxford University Press.

Schubert, S., L. Caviola and N. Faber (2019) 'The Psychology of Existential Risk: Moral Judgments about Human Extinction', *Scientific Reports*, *9*(1), 1-8.

Schwitzgebel, E. and J. Barandes (2024) 'Almost Everything You Do Causes Almost Everything (under certain not wholly implausible assumptions); or Infinite Puppetry', in Eric Schwitzgebel, *The Weirdness of the World*. Princeton: Princeton University Press.

Sidgwick, H. (1907/1874) *The Methods of Ethics*, seventh edition. London: Macmillan.

Smith, Q. (2003) 'Moral Realism and Infinite Space Imply Moral Nihilism', in H. Dyke (ed.), *Time and Ethics: Essays at the Intersection* (Vol. 14). Dordrecht: Springer Science & Business Media.

Tarsney, C. and Wilkinson, H. (2022) 'Longtermism in an Infinite World' *GPI Working Paper* 14-2022. Available at https://globalprioritiesinstitute.org/wp-content/uploads/Working-paper-Christian-J.-Tarsney-Hayden-Wilkinson-Longtermism-in-an-Infinite-World.pdf

Tegmark, M. (2003) 'Parallel Universes', *Scientific American*, *288*(5), 40-51.

Vallentyne, P. and S. Kagan. (1997) 'Infinite Value and Finitely Additive Value Theory', *The Journal of Philosophy* 94 (1), 5–26.

Vilenkin, A. (1983) 'Birth of Inflationary Universes', *Physical Review D*, *27*(12), 2848.

Wallace, D. (2012) *The Emergent Multiverse: Quantum Theory According to the Everett Interpretation*. Oxford: Oxford University Press.

Wilkinson, H. (2021) 'Infinite Aggregation: Expanded Addition', *Philosophical Studies*, 178(6), 1917-1949.

Williams, B. (1981) *Moral Luck: Philosophical Papers 1973-1980*. Cambridge: Cambridge University Press.

Wolf, S. (2010) *Meaning in Life and Why it Matters*. Princeton: Princeton University Press.

Zame, W. (2007) 'Can Intergenerational Equity be Operationalized?' *Theoretical Economics*, 2(2), 187–202.