Repetition and Value in an Infinite Universe

Eric Schwitzgebel
Department of Philosophy
University of California, Riverside
Riverside, CA 92521

May 1, 2023

Repetition and Value in an Infinite Universe

Abstract: On standard physical theory, it's plausible that the universe is infinite and contains infinitely many near-duplicates of you and everything you love. It is also plausible that most of your actions have infinitely many positive and negative effects. This ruins versions of decision theory that rely on summing up all the consequences of one's actions. It opens questions such as: Should we care whether our actions have infinitely many consequences? Would it be better if the cosmos were finite or infinite? Is a cosmos in which everything happens twice, or infinitely many times, twice as good as, or infinitely better than, a cosmos in which everything happens only once? I recommend celebrating the possibility of an infinitely repeating cosmos in which most of our actions have endless effects.

Keywords: cosmology, decision theory, infinity, Nietzsche, recurrence, value summation, value theory

Word Count: about 7700 words

Repetition and Value in an Infinite Universe

Nietzsche writes:

> *The greatest weight.* – What, if some day or night a demon were to steal after you into your loneliest loneliness and say to you: "This life as you now live it and have lived it, you will have to live once more and innumerable times more; and there will be nothing new in it, but every pain and every joy and every thought and sign and everything unutterably small or great in your life will have to return to you, all in the same succession and sequence – even this spider and this moonlight between the trees, and even this moment and I myself. The eternal hourglass of existence is turned upside down again and again, and you with it, speck of dust!"
>
> Would you not throw yourself down and gnash your teeth and curse the demon who spoke thus? Or have you once experienced a tremendous moment when you would have answered him: "You are a god and never have I heard anything more divine." If this thought gained possession of you, it would change you as you are or perhaps crush you. The question in each and every thing, "Do you desire this once more and innumerable times more?" would lie upon your actions as the greatest weight. Or how well disposed would you have to become to yourself and to life *to crave nothing more fervently* than this ultimate eternal confirmation and seal? (Nietzsche 1882/1974, §341, p. 273-274).

There's a respect in which the demon's cosmology is plausible. The universe might well be infinite, containing infinitely many near-duplicates of you, and your actions might have infinite weight, chaotically rippling through an unending future. I'll suggest that granting even the

slenderest credence to such an infinitary cosmology ruins approaches to decision making that require comparing total expected consequences. Nevertheless, we can have reasonable axiological preferences about the general size and structure of the cosmos. I recommend celebrating the possibility of an infinitely repeating cosmos in which our actions have endless effects.

1. The Universe Is Plausibly Infinite, with Infinitely Many Near-Duplicates of You

The limits of what we can *see* are hardly likely to be the limits of what there *is*. It would be quite the un-Copernican coincidence if we happened to be in the exact center of things, with a wall of nothing precisely at the spherical rim of the 46.5 billion light-year range of our telescopes.

Scientific cosmologists commonly think that the universe is in fact infinite (Vilenkin 2006; Tegmark 2009; Linde 2015/2017). This is of course speculative: The infinitude of the universe does not follow straightaway from standard physical theory. However, it is probably the most natural extension of standard physical theory. If the universe is not infinite, it must either have some sort of *edge,* or it must have a *closed topology*. There is no evidence of an edge, nor any widely accepted theory that implies the existence of an edge, and the existence of an edge would require awkward asymmetries and complexities. For example, there would need to be either new and otherwise currently unmotivated physical laws concerning what happens to particles approaching the edge. A closed topology is a more theoretically elegant possibility. The simplest closed topology would be a roughly constant positive curvature of space at very large scale, so that space wraps around upon itself in something like the manner that the surface of a sphere wraps around itself. However, current estimates of large-scale topology suggest that

the portion of the universe we can see is topologically flat (Planck Collaboration 2014). Spatial infinitude appears to be the most straightforward interpretation of the evidence currently available to us.

For similar reasons, it's plausible that the universe endures infinitely. No evidence suggests a temporal edge in the future or motivates the postulation of laws that would govern particles approaching the temporal edge. Nor is there evidence that we exist in a closed temporal loop.

Heat death, of course, follows from standard physical theory. Eventually, the currently observable portion of the universe will settle into a cool, high entropy state. But nothing suggests that time stops at heat death. Instead, standard physical theory suggests that, post-heat-death, physical particles (or waves or fields) will continue to bounce about in the darkness. If particles do continue to exist infinitely into the future, then by chance they will sometimes enter low probability configurations. Seventeen particles will occasionally converge on the same spot – or seventeen million will. There appears to be no in principle limit to the size or structural complexity of such chance fluctuations: If we wait long enough, eventually a molecule-for-molecule near-duplicate of the whole galaxy will arise from the chaos, by pure chance (Boltzmann 1897; Carroll 2021). Wait long enough, and eventually you'll have as many chance-generated galaxies, of diverse form, as you dare to hope for.

As has often been noticed, small chance fluctuations are much likelier than large ones, so brain-sized fluctuations are likelier than galactic-sized fluctuations. This gives rise to the famous "Boltzmann brains" problem: How do you know you aren't a bare brain amid post-heat-death chaos, if there are infinitely many such entities (Carroll 2021; Kotzen 2021)? I won't address the issue here, except to remark that the problem disappears as long as some mechanism generates

new galaxies with new intelligent life sufficiently quickly to eclipse the number of new Boltzmann brains – for example, through the occurrence of new cosmic inflations (De Simone, Guth, Linde, Noorbala, Salem, and Vilenkin 2010).

On the picture developed so far, one way or another, whether through Boltzmannian chance or instead through some other process, such as new inflations, eventually there will be infinitely many future galaxies, presumably instantiating, infinitely often, every possible configuration, however unlikely, that does not have strictly zero probability. If we further assume, as seems reasonable, that the current configuration of our galaxy – perhaps specified within some error tolerance, such as a trillionth of the radius of a proton for every constituent particle – is not a zero-probability event, then infinitely many future near-duplicates of us will live lives qualitatively indistinguishable from our own.

In other words, Nietzsche's demon was right. It's at least *plausible* that the universe is infinite and contains infinitely many near-duplicates of you and all your friends.

2. But Are These Near-Duplicates <u>You</u>?

Nietzsche's demon says one thing that might not fit neatly with the picture so far: that *you* will live innumerable times more. Here we might stop short. Would such future duplicates really be *you?*

I'm inclined to depart from Nietzsche on this point. If a molecule-for-molecule duplicate of you were created, say, on a distant planet, we might not want to say that that person is really you – especially if you continue to exist right here on Earth. Philosophers discussing personal identity typically reject the view that interplanetary duplicates are numerically identical to the people they resemble even if they are qualitatively indistinguishable (Parfit 1984; Kind 2015).

By analogy, it seems reasonable to suppose that a near-duplicate of you in the distant future, especially if it arose by mere chance with no special causal connection to you, would not actually be you yourself. If there's a difference between anticipating that *someone very much like you* will have experiences practically indistinguishable from the experiences you are now having and anticipating that *you yourself* will have these experiences again, the former might be the more correct conceptualization.

3. Almost Everything You Do Causes Almost Everything

If the cosmology articulated so far is correct, then it's probably also true that almost everything you do has effects that ripple infinitely into the future. These effects will be extremely widespread and various.

Suppose you raise your hand now. By doing so, you disturb the trajectory of a huge number of particles: nitrogen and oxygen molecules in the air that would otherwise have taken different paths, photons streaming through the window that now reflect off your thumb instead of striking the wall behind you, volatile organic compounds near the surface of your skin that now wander off in different directions than they otherwise would have. Those disturbed particles then disturb other particles, which disturb other particles, in an ever widening ripple. The effects of this ripple are not confined to Earth. If you're near an open window, for example, a fraction of the photons reflected off your hand will shoot up through the atmosphere into interstellar space, where they will journey until they interact with something – a distant star or planet or piece of dust, for example, which will then behave slightly differently than it otherwise would have, continuing the ripple of effects.

A photon reflected off your hand, let's suppose, perturbs a system which now emits a different photon than it would otherwise have emitted, and that photon shoots out into interstellar space. The photon is absorbed by a black hole, ever so slightly increasing the mass of the black hole and thus ever so slightly changing the trajectory of other photons passing near the black hole but not absorbed by it. Those photons will then perturb other systems differently than they otherwise would have, and so on, deep into the post-heat-death future. Eventually, this ripple will enter a future galaxy. One of these rippling photons, which would not otherwise have been exactly when and where it is, with exactly the wavelength it possesses, will eventually hit the detection surface of a telescope, adding just enough energy to cross a threshold that triggers an alert to a waiting scientist. As a result, the scientist will publish a paper and win a prize, changing her life. Different babies are then born than would otherwise have been born. Different life plans are enacted than would otherwise have been enacted, different poems written, different companies founded, different wars started, different peaces concluded. Although such an outcome has only an extremely tiny probability in any smallish region of future spacetime – multiplying the minuscule upon the minuscule upon the minuscule again – it is presumably not *zero* probability. Eventually a galaxy will be influenced in exactly that way by one of the photons from the ripple leaving your hand. The duration required might make a googolplex-to-the-googolplex-to-the-googolplex-to-the-googolplex years seem like the briefest flash. No problem! We have, after all, infinite time to wait.

Because you raised your hand a minute ago, X happened, and then Y happened, and then Z happened, and then eventually your ripple causes a radioastronomer to win a prize – or any other non-zero-probability event type that you care to name. We can say your action *caused* the scientist to win the prize, if we're not too demanding about what counts as a "cause". If you

hadn't raised your hand, it wouldn't have happened – not then and there, to that particular scientist – and there is a continuous chain of causal physical processes from the lifting of your hand to the winning of the prize. In this weak sense of causation, almost everything you do causes almost every type of finitely specifiable, non-zero-probability, non-unique type of physically possible future event.

Call the cosmology of Sections 1 and 3 the *Infinitary Cosmology*. This cosmology is, I hope, physically plausible – the most straightforward extension of current physical theory. In other work (Schwitzgebel and Barandes forthcoming), I have defended its physical plausibility in more detail. The remainder of this article explores the axiological consequences.

4. Some Problems with Infinite Expected Values

The Infinitary Cosmology might seem to imbue our actions with a potentially troubling weight: By raising your arm right now, you will cause infinitely many future deaths, due to the causal ripple emanating from that action. Superficially, that might seem like a good reason not to raise your arm. But of course, if the Infinitary Cosmology is correct, holding your hand motionless will also cause infinitely many (different) future deaths. Of course, both actions will also prevent infinitely many future deaths. What are we to make of this?

For one thing, it seems to ruin some approaches to action evaluation. According to some standard versions of consequentialist ethics and ordinary decision theory, the goodness or badness of your actions depends on their overall consequences: the sum total of the positive consequences minus the sum total of the negative consequences. If the Infinitary Cosmology is correct, the sum total value of almost all of your actions will be $\infty + -\infty$, a sum which is normally considered to be mathematically undefined. Suppose you are considering two possible actions

with short-term expected values *m* and *n*. Suppose, further, that *m* is intuitively much larger than *n.* Maybe Action 1, with short-term expected value *m,* is donating a large sum of money to a worthwhile charity, while Action 2, with short-term expected value *n,* is setting fire to that money to burn down the house of a neighbor with an annoying dog. The Infinitary Cosmology breaks the mathematical apparatus for comparing the overall value of those actions: The total expected value of Action 1 will be $m + \infty + -\infty$, while the total expected value of Action 2 will be $n + \infty + -\infty$. Both values are undefined.

An Optimist might try to escape the problem thus: Suppose that overall in the universe, at large enough spatiotemporal scales, the good outweighs the bad. We can now consider the relative values of Action 1 and Action 2 by dividing them into three components: the short-term effects (*m* and *n*, respectively), the medium-term effects *k* – the effects through, say, the heat death of our region of the universe – and the infinitary effects ($+\infty$, by stipulation). Stipulate that *k* is unknown but expected to be finite and similar for Actions 1 and 2. The expected value of Action 1 is thus $m + k + \infty$. The expected value of Action 2 is $n + k + \infty$. These values are not undefined; so that particular problem is avoided. The values are, however, equal: simple positive infinitude in both cases. As the saying goes, infinity plus one just equals infinity. A parallel Pessimistic solution – assuming that at large enough time scales the bad outweighs the good – runs into the same problem, only with negative infinitude.

Perhaps a solution is available for someone who thinks that at large enough time scales the good will exactly match the bad, so that we can compare $m + k + 0$ to $n + k + 0$? Positive and negative will balance exactly, as if on a knife's edge. The problem with the Knife's Edge solution is delivering that zero. Even if we assume that the *expected* value of any finite spatiotemporal region is exactly zero, the Law of Large Numbers only establishes that as the

number of finite regions under consideration goes to infinity, their *average value* is very likely to be near zero. The *sums* will not converge upon specific values. If good and bad effects are randomly distributed and do not systematically decrease in absolute value over time, then the relevant series would be $m + a_1 + b_1 + c_1 + d_1 + \ldots$ and $n + a_2 + b_2 + c_2 + d_2 + \ldots$ where each variable after the first can take a different positive or negative value and where there is no finite limit to the value of positive or negative runs within the series. These are seemingly the very archetype of poorly behaved divergent series with incalculable sums (even by clever tools like Cesàro summation). Thus, mathematically definable sums still elude us.

Perhaps the advocate of Knife's Edge reasoning can shift to evaluating consequences based on their average effects, taking the limit of the *ratio* of good to bad effects as the number of effects goes to infinity? This will leave them in the same place as the Optimist or Pessimist, indifferent between Actions 1 and 2, since the limit will be the same in both cases – zero – washing out the finite values of m vs n. Non-zero limits will similarly result in indifference.

Could we then compare $m + a_1 + b_1 + c_1 + d_1 + \ldots$ and $n + a_2 + b_2 + c_2 + d_2 + \ldots$ by treating all terms beyond the first as identical in expectation, then subtracting from both sides? No, we can't do this either, for reasons well-known to those familiar with paradoxes of infinitude. When applied to infinite series, intuitive principles of grouping and linear transformation lead to absurdities such as the Hilbert's Hotel paradox and the conclusion that $1 + 2 + 3 + 4 + \ldots = -1/12$ (Dodds 2018). For good reason, mathematicians evaluate infinite series by looking at their limits rather than by ordinary rules of linear transformation. If infinitude plus one is no different from infinitude, so also is infinitude plus m no different from infinitude plus n.

Could we appeal to dominance reasoning? According to dominance reasoning, if Action A has better results than Action B regardless of what else happens, Action A should be chosen. This might justify the choice of, say, a bet than pays \$1000 plus $\$2^n$ (where n is the number of times a fair coin lands heads) over a bet that just pays $\$2^n$ despite the fact that both bets have infinite expectation (Hajek and Nover 2006; Easwaran 2021; Wilkinson 2021). However, dominance reasoning doesn't apply to the present case, since it is not true that Action 1 will have better results than Action 2 no matter what else happens.

All of this generates a dilemma for believers in the Infinitary Cosmology who hope to evaluate actions by their total expected value. Either accept the conclusion that there is no difference in total expected value between donating to charity and burning down your neighbor's house (the solution of the Optimist, Pessimist, or the theorist who prefers ratios rather than sums), or accept that there is no mathematically definable total expected value for any action, rendering proper evaluation impossible.[1]

The solution, I suggest, is not to evaluate actions based on their total expected value over the lifetime of the cosmos! We must have some sort of discounting with spatiotemporal distance, or some limitation of the range of consequences we are willing to consider, or some other policy to expunge the infinitudes from our equations. Unfortunately – as Bostrom (2011) persuasively argues – no such solution is likely to be entirely elegant and intuitive from a formal point of view: Fancy mathematics doesn't handle all the plausible cases, and various discounting

[1] Wilkinson 2021 might seem to offer a decision-theoretic solution, but in fact their preferred resolution of the analogous case ("Writing or Netflix") finds a comparison of the outcomes to diverge to *either* $+\infty$ or $-\infty$ with equal probability, and thus no decision-theoretical basis to choose charity over house-burning. See also Lenman 2000 and Greaves 2016 on "cluelessness" and Chappell 2001 on "option ranges".

regimes appear to generate unintuitive consequences. So much the worse, then, for an elegant, intuitive, complete, formal model of infinitary decision making.

The infinite expectation problem is robust in two ways. First, it affects not only simple consequentialists. After all, you needn't be a simple consequentialist to think that long-term expected outcomes matter. Virtually everyone thinks that long-term expected outcomes matter *somewhat*. As long as they matter enough that an infinitely positive long-term outcome would be relevant to your evaluation of an action, you risk being caught by this problem.

Second, the problem affects even people who regard the Infinitary Cosmology as unlikely. Even if you are 99.99% certain that the Infinitary Cosmology is false, your remaining 0.01% credence in the Infinitary Cosmology will destroy your expected value calculations if you don't somehow sequester the infinitudes. Suppose you're 99.99% sure that your action will have the value *k*, while allowing a 0.01% chance that its value will be $\infty + -\infty$. Now apply the expected value formula in the standard way. Unfortunately, $.9999 * k + .0001 * (\infty + -\infty)$ is just as undefined as $\infty + -\infty$ itself. Similarly, $.9999 * k + .0001 * \infty$ is simply ∞. As soon as you let those infinitudes influence your decision, you risk falling back into the dilemma.

5. Should We Care Whether Our Ripples Continue?

I have suggested that the Infinitary Cosmology is plausible and that investing even a tiny credence in the Infinitary Cosmology ruins any attempt to evaluate actions by their total expected consequences. It might seem to follow that we should *ignore* the fact that our actions plausibly have infinitely many good and bad consequences. However, I don't think we should ignore that fact.

Consider two alternatives to the Infinitary Cosmology:

The Small Cosmology. Spatially, the cosmos is not much bigger than what we can see, and temporally, it is also limited. After heat death, all existence ceases. Perhaps time itself comes to a stop. The ripples of your actions, of course, also cease.

The Erasure Cosmology. The cosmos endures infinitely, but at some point every rippling consequence of your actions is stopped. Perhaps our region concludes in a Big Crunch, which launches a new Big Bang. However, this Big Bang happens exactly as it would have happened regardless of any action of yours. Whether your hand goes up or stays down, the new Big Bang proceeds exactly the same either way. No trace of your actions remains post-Crunch.

It is not, I think, unreasonable to evaluate these three cosmologies differently. In the Infinitary Cosmology, everything you do has consequences, good and bad, infinitely into the future. You cannot do anything now in the reasonable expectation of favorably changing the overall balance of good versus bad, but still, the infinite future transpires differently than it otherwise would, and eventually your actions will have caused virtually every non-zero probability event infinitely many times. Whether this is appealing might depend on your personality or values.

Some people might hope that their influence on the universe will eventually cease, though the universe continues. They might like the idea of departing the world without a trace, so that after some point the universe continues exactly as it would have continued had they never existed. They'd like to walk through the world as one might walk through a forest, leaving only footprints that fade away, influencing nothing, so that eventually others might walk through the forest entirely unaffected by their earlier passage. For the sake of this analogy, let's ignore the fact that, realistically, footprints will influence bugs, which will influence birds, etc. Such a person might prefer the Erasure Cosmology.

Others might like the idea that the entire cosmos ceases, their own traces along with everything else, thus preferring the Small Cosmology. In the next section, I'll discuss the Small Cosmology more, but here I just note that it is very different from the Erasure Cosmology.

Still others might rather enjoy the idea that their effects will ripple forward infinitely through time. You raise your hand, starting a ripple that eventually in the far, far future causes some astronomer to win a prize. If not for this action, she would not have won that prize. Perhaps this is weirdly wonderful, weaving you more deeply into the cosmos. There might also be an infinite past – some cause of our own Big Bang, tracing back to prior influences, influences rippling from some long-ago intelligent entity thinking cosmologically about their possible influence on the very distant future. Your actions would carry the traces of their actions, much as the future radioastronomer's actions carry the traces of your actions. We walk through the forest, and our traces do not vanish. The forest is not ruined, but it continues differently. Our footprints redirect the bugs, who redirect the birds, and so on, and a thousand years later there's a robin on a willow singing differently.

Is there any ethical reason why we should hope that our ripples continue? I doubt it. No overall good would seem predictably to come of it, nor does it satisfy any obvious imperative. If anything, it might be admirably modest to hope one's ripples cease. Nor is there any clear prudential reason to hope our ripples continue, some straightforward way in which our own lives go better, if the Infinitary Cosmology is true.

Aesthetically, one might have preferences. There's perhaps something beautiful in a picture of the cosmos in which the ripples of our actions continue infinitely, intertwining with infinitely many future lives, in every combination, over and over. But perhaps there's also something beautiful in a picture on which the ripples all eventually end, while everything else

continues. Perhaps there's even something beautiful or sublime in the thought of a final end of everything.

Axiologically, however, I think we should not be indifferent about the size of the universe. An infinite universe is arguably overall better than an finite one, as I will now discuss.

6. Optimism, Pessimism, and Hopes for the Size of the Cosmos

The Optimist, let's say, holds that, at large enough spatiotemporal scales, the good outweighs the bad. Put differently, as the size of a spherical spatiotemporal patch grows extremely large it becomes extremely likely that the good outweighs the bad. Optimism would be defensible on hedonic grounds if the following is plausible: At large enough scales, the total amount of pleasure will almost certainly outweigh the total amount of pain, among whatever inhabitants occupy the region. The Pessimist holds the opposite: At large enough spatiotemporal scales, the bad outweighs the good – perhaps, again, on hedonic grounds, if the pain outweighs the pleasure. A Knife's-Edge theorist expects a balance.

I see no good hedonic defense of Optimism. Suffering is widespread and might easily balance or outweigh pleasure. I prefer to defend Optimism on eudaimonic grounds: Flourishing lives are valuable, and flourishing lives are virtually guaranteed to occur in sufficiently large spatiotemporal regions.

Imagine a distant planet – one on the far side of the galaxy, blocked by the galactic core, a planet we will never interact with. What ought we hope this planet is like, independent of its relationship to us? Ought we hope that it's a sterile rock? Or would it be better for the planet to host some sort of life? If the planet hosts some sort of life, would it be best if that life is only simple, microbial life, or would complex life be better – plants, animals, and fungi, savannahs

and rainforests and teeming reefs? If it hosts complex life, would it be better if nothing rises to the level of human-like intelligence? Or ought we hope for societies, with families and love and disappointment and anger, poetry and philosophy, art and athletics and politics, triumphs and disasters, heroism and cruelty – the whole package of what is sometimes wonderful and sometimes awful about human existence?

A Pessimist might say the sterile rock is best – or rather, least bad – presumably because it has the least suffering and vice. But I suspect the majority of readers will disagree with the Pessimist. Most, I suspect, will believe, as I do, that complex life is better than simple life, which is better than sterility, and that what's most worth hoping for is the full suite of love, poetry, philosophy, science, art, and so on. The galaxy overall is better – more awesome, wondrous, and valuable – if it contains a distant planet rich with complex life, a bright spot of importance. If something were to wipe it out or prevent it from starting, that would be a shame and a loss. On this way of thinking, Earth too is a bright spot. As a general matter – perhaps with some miserable exceptions – complex life is not so terrible that nonexistence would be better. The Pessimist is missing something.

What form, then, should we hope the cosmos takes?

A benevolent Pessimist might hope for the Small Cosmology, on the principle that the Small Cosmology contains only finitely much badness, and finite badness is better than infinite badness. (A spiteful Pessimist might hope for infinite badness.) Presumably nothingness would have been even better. A less simple Pessimism might hold that the observable portion of the universe is already infinitely bad. This might entail indifference about the existence or nonexistence of additional regions, depending on whether the infinitudes can be compared. Another less simple Pessimism might suspect that the observable portion of the universe is worse

than the average spatiotemporal region and so hope for enough additional material to bring the average badness of the cosmos to a more acceptable level. Still other forms of Pessimism are of course conceivable, with some creative thinking.

But we are, I hope, Optimists. Some Optimists might hold that the observable portion of the universe is infinitely good. If so, they might conclude that a larger cosmos would not be better unless they're ready to weigh the infinitudes differently. More moderately and plausibly, the observable portion of the universe might be only finitely good. Call this view Muted Optimism.

Here's one argument for Muted Optimism. Suppose you agree that if a human life involves too much suffering, it is typically not worth living. By analogy, it seems plausible that if the observable portion of the universe contained too much suffering, it would be better if it didn't exist. We needn't be hedonists to accept this idea. Contra hedonism, flourishing life might be overall good despite containing more suffering than pleasure. It just might not be *so* good that there isn't some amount of suffering that would make the combined package worse than nothing. But if flourishing were infinitely good, then no amount of suffering could outweigh it (though infinite suffering might create a $\infty + -\infty$ situation). Therefore, large finite regions are good but not infinitely good.

Muted Optimism suggests that an infinite cosmos would be better than the Small Cosmos. It seems, after all, that more goodness is better than less goodness, and infinite goodness seems best. As with Pessimism, however, the axiology needn't be quite so simple. For example, one might hold that too much of a good thing is bad. Or one might suspect that the observable portion of the universe is much better than could reasonably be expected from a typical region and that adding more regions would objectionably dilute average goodness. Or

one might simply think it would be stupendously awesome if the cosmos were some particular finite size – shaped like a giant jelly donut, perhaps, with red galaxies in the middle and lots of organic sugars along the edges.

Or one might mount the Repetition Objection, to which I will now turn.

7. Repetition and Value in an Infinite Cosmos

Consider a particular version of the Erasure Cosmology. There's a Big Bang, things exist for a while, and then there's a Big Crunch. Suppose that what happens next is an exact repetition of the first Bang-existence-Crunch. You, or rather a duplicate of you, lives exactly the same life, having exactly the same experiences, seeing exactly the same moonlight between the trees and having exactly the same thoughts about that moonlight, as envisioned by Nietzsche, all over again. And then it happens again and again, infinitely often. Call this Repetitive Erasure.

Now contrast this picture with the same cosmos, except that after the Crunch nothing exists. Call this cosmos Once and Done. Finally, contrast these two possibilities with a third, in which there is exactly one repetition: Twice and Done. (If you're inclined toward metaphysical quibbles about the identity of indiscernibles, let's imagine that each Bang and Crunch has some unique tag.)

How might we compare the values of Once and Done, Twice and Done, and Repetitive Erasure? Four simple possibilities include:

Equal Value. Once and Done, Twice and Done, and Repetitive Erasure are all equally good. There's no point in repeating the same events more than once. But neither is anything lost by repetition.

Linear Value. If Once and Done has value *x*, then Twice and Done has value 2*x*, and Repetitive Erasure has infinite value. The value of one run-through is not diminished by the existence of another earlier or later run-through, and the values sum.

Diminishing Returns. If Once and Done has value *x*, then Twice and Done has a value greater than *x* but less than 2*x*. Repetitive Erasure might have either finite or infinite value, depending on whether the returns converge toward a limit. A second run-through is good, but two run-throughs are not twice as good as a single run-through: Although it's not the case that there's *no* point in God's hitting the replay button, so to speak, there's less value in running things twice.

Loss of Value. If Once and Done has value *x*, then Twice and Done has a value less than *x*, and Repetitive Erasure is worse, perhaps even infinitely bad.

If Equal Value or Loss of Value is true, then Muted Optimism shouldn't lead to preference for the infinitude of Repetitive Erasure over the finitude of Once and Done. If we further assume that in an infinite cosmos, the repetition (within some error tolerance) of any finite region is inevitable, then the argument appears to generalize. This is the Repetition Objection. Some positively-valenced existence is good, but after a point, more of the same is not better (e.g., Bramble 2016).

In ordinary cases, uniqueness or rarity can add to a thing's value. One copy of the Mona Lisa is extremely valuable. If there were two Mona Lisas, presumably each would be less valuable, and if there were a billion Mona Lisas no one of them would presumably be worth much at all. The question is whether this holds at a cosmic scale. Might this only be market thinking, reflecting our habit of valuing things in terms of how much we would pay in conditions

of scarcity? Or is there in fact something truly precious in uniqueness? (For discussion, see Lemos 2010; Chappell 2011; Bradford forthcoming.)

Perhaps there is something beautiful, or right, or fitting, in things happening only once, in a finite universe, and then ceasing. Is it good that you are the only version of you who will ever exist, so to speak – that after you have lived and died there will never again be anyone quite like you? Is it good that the cosmos contains only a single Confucius and only a single Great Barrier Reef, no duplicates of which will ever exist? Things will burn out, never to return. There's a romantic pull to this idea.

Against The Repetition Objection to the simple Muted Optimist's preference for an infinite universe, I offer the Goldfish Argument (see also Schwitzgebel 2019, ch. 44).

According to popular belief (not in fact true), goldfish have a memory of only thirty seconds. Imagine, then, a goldfish swimming clockwise around a ring-shaped pool, completing each circuit in two minutes. Every two minutes it encounters the same reeds, the same stones, and the same counterclockwise-swimming goldfish it saw in the same place two minutes before, and each time it experiences all of these as new. The goldfish is happy with its existence: "Howdy, stranger, what a pleasure to meet you!" it says to the counterclockwise-swimming fish it meets afresh every minute. To tighten the analogy with the Repetitive Erasure cosmology, let's stipulate that each time around this goldfish sees and does and thinks and experiences exactly the same things.

Now stop the goldfish mid-swim and explain the situation. The goldfish will not say, "oh, I guess there's no point in my going around again." The goldfish will want to continue its happy little existence, and rightly so. It still wants to see and enjoy what's around the next bend. Moment to moment it is having good experiences. You harm and disappoint the goldfish by

stopping its experiences, as long as each experience is, locally, good – even if they have all happened before innumerably many times. This is true whether we catch the goldfish after its first swim around, after its second, or after its googolplex-to-the-googolplexth. It's better to let the fish swim on. If the analogy holds at cosmic scales, then Equal Value and Loss of Value must be false. Maybe, though, there's still something attractive about uniqueness, some truth in it that isn't simply inappropriate market-style thinking? I see no need to deny that there really is something special about the first time. Let's grant that it's possible that the first go-round is somehow made less valuable by later go-rounds. As long as the harm done by stopping the goldfish (by denying future goods) exceeds the harm done by letting the goldfish continue (by reducing the rarity of past goods), then Diminishing Returns is the correct view. If we further assume that the added value does not continually shrink in a way that approaches zero, then the view we should embrace is one on which Repetitive Erasure would have infinite value.

This thinking appears to extend to the Infinitary Cosmology. Duplicates of you, and me, and all Earth, and the whole Milky Way will repeat over and over, infinitely. Each repetition adds some positive value to the cosmos, and in sum the value is infinite.

8. Replying to Nietzsche's Demon.

The Muted Optimist might reply to Nietzsche's demon thus: "Demon, though I doubt we should call these future duplicates *me* strictly speaking, your cosmology is plausible. When you say these duplicates will live exactly as I have lived, you leave out part of the story. Infinitely many will do so, but another infinitude will also live every other life a duplicate of me could possibly lead. For every choice I made or will make, future counterparts of me will make different choices, some better, some worse. Some will be carpenters, some itinerant street

musicians, and some Vice Presidents of the Union of One-Armed Spelunkers. Every possible accident will befall them. Infinitely many will drown at age six. Infinitely many will win the state lottery in a future duplicate of Poland. Infinitely many will fly up on a freakish gust of wind, then settle gently back down.

"You imply that my actions have immense weight, echoing through these future versions of me. I agree. Infinitely many of these future duplicates will do the things they do in part because of what I do now, speaking to you. Due to the rippling effects of my actions, if we hadn't had this conversation, infinitely many particular future versions of myself would not have existed or would have acted differently than they did. An infinite number of my future counterparts are in this way tied to me – just as infinitely many other people are also influenced by my actions, which resonate unendingly through the cosmos. This gives my actions, in a sense, infinite weight. But I do not attempt to choose based on those infinite future consequences.

"We are all in the same position. Virtually every thinking being who has ever existed will repeat with infinite variation and have infinitely various incalculable effects upon the future. We are tied together in an endless web of positive value. It is a gloriously weird and awesome vision of reality, Demon. Let's hope that you are right."[2]

[2] For helpful discussion, thanks to Jacob Barandes, Ben Bramble, Sean Carroll, Richard Chappell, Kenny Easwaran, Stephen Hetherington, Linus Huang, Eric Steinhart, and commenters on relevant posts on my blog and social media pages.

References:

Boltzmann, Ludwig (1897). Zu Hrn. Zermelo's Abhandlung "Ueber die mechanische Erklärung irreversibler Vorgänge". *Annalen der Physik, 296,* 392-398.

Bostrom, Nick (2011). Infinite ethics. *Analysis and Metaphysics, 10,* 9-59.

Bradford, Gwen (forthcoming). Uniqueness, intrinsic value, and reasons. *Journal of*

Bramble, Ben (2016). A new defense of hedonism about well-being. *Ergo, 3,* #4.

Carroll, Sean (2021). Why Boltzmann Brains are bad. In S. Dasgupta, R. Dotan, and B. Weslake, eds., *Current controversies in philosophy of science.* Routledge.

Chappell, Richard Yetter (2011). *Value holism.* Unpublished manuscript at https://philpapers.org/archive/CHAVH.pdf

Chappell, Sophie-Grace (2001). Option ranges. *Journal of Applied Philosophy, 18,* 107-118.

De Simone, Andrea, Alan H. Guth, Andrei Linde, Mahdiyar Noorbala, Michael P. Salem, and Alexander Vilenkin (2010). Boltzmann brains and the scale-factor cutoff measure of the multiverse. *Physical Review D,* 82: 063520.

Dodds, Mark (2018). The Ramanujan summation: $1 + 2 + 3 + \ldots + \infty = -1/12$? Blog post at *Cantor's Paradise* (Sep 2).

Easwaran, Kenny (2021). A new method for value aggregation. *Proceedings of the Aristotelian Society, 121,* 299-326.

Greaves, Hilary (2016). Cluelessness. *Proceedings of the Aristotelian Society, 116,* 311-339.

Hájek, Alan, and Harris Nover (2006). Perplexing expectations. *Mind, 115,* 703-720.

Kind, Amy (2015). *Persons and personal identity.* Polity.

Kotzen, Matthew. (2021). What follows from the possibility of Boltzmann Brains? In S. Dasgupta, R. Dotan, and B. Weslake, eds., *Current controversies in philosophy of science.* Routledge.

Lenman, James (2000). Consequentialism and cluelessness. *Philosophy & Public Affairs, 29,* 342-370.

Lemos, Noah (2010). Summation, variety, and indeterminate value. *Ethical Theory and Moral Practice, 13,* 33-44.

Linde, Andrei (2015/2017). A brief history of the multiverse. *ArXiv* 1512.01203v3.

Nietzsche, Friedrich (1882/1974). *The gay science*, trans. W. Kaufman. Random House.

Parfit, Derek (1984). *Reasons and persons.* Oxford University Press.

Planck Collaboration (2014). Planck 2013 results: XVI cosmological parameters. *Astronomy & Astrophysics, 571*: A16.

Schwitzgebel, Eric (2015). Against the "still here" reply to the Boltzmann brains problem. Blog post at *The Splintered Mind* (Oct. 1).

Schwitzgebel, Eric (2019). *A theory of jerks and other philosophical misadventures.* MIT Press.

Schwitzgebel, Eric, and Jacob Barandes (forthcoming). Almost everything you do causes almost everything (under certain not wholly implausible assumptions); or, infinite puppetry. In E. Schwitzgebel, *The weirdness of the world.* Princeton University Press.

Tegmark, Max (2009). The multiverse hierarchy. *arXiv* 0905.1283.

Vilenkin, Alex (2006). *Many worlds in one.* Hill and Wang.

Wilkinson, Hayden (2021). Infinite aggregation: Expanded addition. *Philosophical Studies, 178,* 1917-1949.