



Was evolution worth it?

Guy Kahane^{1,2} 

Accepted: 30 September 2022
© The Author(s) 2022

Abstract

The evolutionary process involved the suffering of quadrillions of sentient beings over millions of years. I argue that when we take this into account, then it is likely that when the first humans appeared, the world was already at an enormous axiological deficit, and that even on favorable assumptions about humanity, it is doubtful that we have overturned this deficit or ever will. Even if there's no such deficit or we can overturn it, it remains the case that everything of value associated with humanity was made possible by our evolutionary history and all that animal suffering. It can seem indecent to regard all that past suffering as having been worth it simply because it was a causal precondition for our existence. But when we consider the realistic alternatives to the way evolution in fact unfolded, there is nevertheless a conditional case for regarding past sentient suffering as a kind of necessary evil.

Many of us feel awe when we reflect on the blind evolutionary process that gave rise to all living things as well as, eventually, to *Homo sapiens*, making possible things such as Bach's music or quantum mechanics. Darwin thought there is grandeur in a view of life where "endless forms most beautiful and most wonderful" evolve from "so simple a beginning",¹ and for Dennett, the evolutionary "tree of life" is something magnificent, even sacred.² And precisely because Darwin showed that humans have no privileged place in creation, we might feel, with Dawkins, that "[w]e have

¹ Darwin (1859), 490.

² Dennett (1995), 520.

✉ Guy Kahane
guy.kahane@philosophy.ox.ac.uk

¹ Professor of Moral Philosophy, University of Oxford, Oxford, United Kingdom

² Tutor and Fellow in Philosophy, Pembroke College, Oxford, United Kingdom

cause to give thanks for our highly improbable existence, and the law-like evolutionary processes that gave rise to it.”³

But evolution should also inspire our horror. Not self-pity about our ‘lowly’ origins but pity for all those quadrillions (perhaps quintillions) of sentient beings who suffered over hundreds of millions of years of relentless competition for scarce resources, a seemingly endless cycle of devouring and beings devoured.⁴ Dawkins elsewhere writes that

During the minute it takes me to compose this sentence, thousands of animals are being eaten alive; others are running for their lives, whimpering with fear; others are being slowly devoured from within by rasping parasites; thousands of all kinds are dying of starvation, thirst and disease.”⁵

This is a claim about the amount of animal suffering in the wild over one minute. Multiply that by the number of minutes in a year, and what you get is an annual hell, something “beyond all decent contemplation”.⁶ And when we turn to evolution, we need to multiply that annual cycle of natural evil hundreds of millions of times over, all the way back to the moment the first sentient being first opened its eyes.

That the pre-human past contains a staggering amount of suffering is hardly news—reflecting on the routines of predation in nature, and on the “carnivorous reptiles of geologic times”, William James wrote that “the deadly horror which an agitated melancholiac feels is *the literally right reaction*”.⁷ But it’s also not something that gets much attention. My aim here is to explore the wider significance of this bleak fact about the terrestrial past.

Bernard Williams wrote that, with its typical focus on agents and their acts, moral philosophy leaves little space for seeing the great horrors of the past as a central concern.⁸ By contrast, Williams portrays Nietzsche as someone who refused

“to forget, not only the existence of suffering, but the fact that suffering was necessary to everything that he and anyone else valued.”⁹

And, Williams writes, since

“the world’s achievements and glories—art, self-understanding, nobility of character—cannot in common honesty be separated from the knowledge of the horrors that have been involved in bringing these things about... there is a ques-

³ Cited in Nagasawa (2018). For simplicity’s sake, I’ll use ‘human’ and ‘Homo sapiens’ interchangeably even though we’re not the only humans who had existed.

⁴ This is an estimate—the actual number may well be even larger. I draw here on an estimate of the number of wild animals currently in existence—see Tomasik, B. (2019).

⁵ Dawkins (1995), 131–132.

⁶ Dawkins (1995), 131.

⁷ James (1929), (my italics).

⁸ Williams (2009a).

⁹ Williams. (2001), xiv.

tion that cannot, Nietzsche supposed, simply be ignored: whether it has all been worth it.”¹⁰

Commenting on this passage, Parfit remarks that

“In asking whether human history has been worth it, we are asking whether the horrors and the suffering have been outweighed, so that human history has been, on the whole, good.”¹¹

This is the question to which pessimists such as Schopenhauer answer ‘No’, because on their view “human existence is on the whole bad”.¹²

But Parfit’s question about overall value is different, I believe, from that posed by Williams’s Nietzsche.¹³ The overall value of human history, or the world as a whole, depends on what of value these spatiotemporal wholes contain, not on the nature of the causal links between value and disvalue. History could be overall bad even if the good in it didn’t depend on the bad. And we can wonder whether humanity’s great achievements were worth it, given the horrors that they required, even if we’re confident that human history is overall good. Moreover, when we ask about the value of wholes, we just need to look at what they actually contain. But to suggest, as Williams does, that suffering was *necessary* for humanity’s greatest achievements, is to make a claim about how things *could* have unfolded. An idea that Williams explores (but doesn’t exactly endorse) is that it might have all been ‘worth it’ if the resulting good was great enough, compared to the bad, and there really was no way to realize this good without the suffering.¹⁴

Parfit and Williams are offering different conditions for it to be the case that ‘it has all been worth it’: one concerned with actual *overall* value—and the rejection of pessimism—and one concerned with *modal* questions about the necessary role of evils. Both conditions, I believe, are needed for a full account of when a span of history can be said to have been worth it. For things to be good on the whole is only a necessary condition for a history to have been worth it, since even an overall positive balance would not ‘be worth it’ if the goods in it (or broadly equivalent or greater goods) could have been realistically realised without those (or equivalent or even worse) evils.¹⁵

¹⁰ Williams. (2001), xiv.

¹¹ Parfit (2011), 607.

¹² *Ibid.*, 607.

¹³ Nietzsche was in fact dismissive of the idea that the world has an overall value, writing that “[t]he total value of the world is unevaluable, consequently philosophical pessimism is among the comical things.” (Nietzsche, 2003, 212).

¹⁴ Williams comes closest to endorsing a version of this in Williams (2009b), though, characteristically, he understands the question in psychological terms: suffering was worth it when we can bear to honestly face it, and we can do so when we see it as playing an essential role in our key projects.

¹⁵ At least in principle, the goods and evils in a course of history could be entirely causally independent; but such evils would even more clearly fail to count as necessary evils.

Parfit and Williams are explicitly concerned with human history.¹⁶ But Parfit also speaks of the value of the world as a whole.¹⁷ And even if human history was net positive, it's hard to see why this would make it all worth it if human history was preceded, and outweighed, by a mind-bogglingly greater amount of animal suffering. To decide whether or not pessimism is true, we must take that suffering into account as well. This will be my focus in the first part of the paper. Mind-boggling as it may be, the suffering associated with evolution cannot establish global pessimism on its own. But this suffering does make it likely that when the first humans appeared, there was already an enormous axiological deficit. And if there was such a deficit, it's far from obvious that we humans have overturned it, or ever will. Arguments for pessimism typically make revisionary claims about human lives, claims that many reject. But the evolutionary case for thinking that pessimism is a serious prospect makes no such revisionary claims and is even compatible with a rosy view of human lives.

Even if we can overturn that deficit, this still wouldn't show that this suffering was worth it in the sense suggested by Williams. And humanity's great achievements, and our very existence, depend not only on past injustice and brutality by humans towards their fellow humans (not to mention other animals), but also on billions of years of 'law-like evolutionary processes' and the suffering they involved. On reflection, it can seem monstrous to think that we should give thanks, with Dawkins, for such a horrific process simply because it led to us. In the second part, however, I'll argue that when we consider the realistic alternatives to the actual past—when we consider whether and how things could have been better, and for whom—and so long as there's still a chance that things will turn out good on the whole, then a conditional case can be made for regarding all that animal suffering as a necessary evil. Although the animal suffering associated with evolution means that there's a serious chance that Parfit's overall value condition *won't* be met, if it *will* be met, it is likely that this past suffering will also meet Williams's Nietzschean modal condition. However, these reflections also suggest the possibility of an even bleaker form of pessimism.¹⁸

¹⁶ Though Parfit mentions animal suffering in passing (Parfit, 2011, 613).

¹⁷ *Ibid.*, 620.

¹⁸ The question of whether past suffering has been worth it can sound similar to the more familiar question of whether the seemingly senseless suffering we see around us can be reconciled with the existence of an omniscient, omnipotent and perfectly good God. And from Darwin onwards, the suffering of wild animals—a 'natural evil' that cannot be traced to the faulty exercise of free will—has often been seen as especially hard to reconcile with God's existence (for an attempt to address such evil, see Murray (2008)). With Nietzsche, Williams and Parfit (and Dennett and Dawkins), I will simply assume here a broadly naturalist framework—in part *because* I doubt there's a good reply to the problem of evil. The problem of evil is supposed to give us a reason to conclude that God doesn't exist. But if we atheists conclude that past animal suffering wasn't worth it, in the sense I have sketched above, this shouldn't change our view about what exists (contra Nagasawa 2018; see Kahane 2021b). It should just change the attitudes we have towards the past, and our own existence. Conversely, if we conclude that evolution *was* worth it, in our sense, this in no way blunts the force of the problem of evil. It would only mean that, given the context of the narrow range of scenarios that were possible within a naturalist framework, scenarios involving massive evolutionary suffering needn't be regretted, all things considered. This in no way entails that it would have been permissible for someone to intentionally bring about such a sequence of events, let alone a perfect supernatural being for whom *all* possibilities are realistic (see Kahane, forthcoming; for the view that God's range of relevant options here may have been far more narrow, see van Inwagen, 2008). I will return to the question of theism at the very end of this paper.

1 Are things bad on the whole?

1.1 The case for evolutionary pessimism

The mass suffering associated with evolution merits our horror. But such badness also has wider axiological significance, by bearing on the value of the larger wholes that contain it. How that suffering bears on larger wholes will depend on what axiology we adopt. To make the discussion manageable, I'll largely assume the following. First, that the value of a whole is the sum of the final value it contains, though I will also briefly consider alternative views. Second, I'll largely understand animal well-being in hedonic terms. Evolutionary history is awful primarily because it contains so much agony—largely physical pain and other negative physical sensations, but in higher animals also emotions such as fear and distress. For short, I'll refer to all that suffering as *evolutionary evil*. Third, I'll assume, with many others, that sentience is a necessary condition for value. This entails that prior to the moment that first sentient creature stirred, value was completely absent from the cosmos.¹⁹

The logic of evolution has been at work for about 3.5 billion years. Over this period, many quadrillions of living things were spawned, and then subjected to predation, parasitism, inter- and intra-species aggression, starvation, and disease. At some point in this process—probably hundreds of million years ago²⁰—sentience first emerged, and from that moment onwards, the evolutionary process inherently involved, not just death and decay, but also rampant suffering. We can start by asking how this evolutionary evil bears on the value of the most immediate larger whole that contains it—that of evolutionary history itself. The question is whether the mass suffering involved in evolution means that what I'll call *Evolutionary Pessimism* (EP) is true: that terrestrial evolutionary history is on the whole bad.

However, throughout that period, many trillions of sentient beings also experienced numerous episodes of pleasure—feeding, mating, basking in the sun... And these positive episodes also add up to a vast quantity of good. Few ask whether the balance of pain and pleasure in nature is positive. But many think of nature in rather rosy terms, an attitude that's hard to square with endorsement of EP. There is also, of course, a pessimist tradition that, with Schopenhauer, just finds it obvious that the natural world contains vastly more suffering than pleasure.²¹ This disagreement cannot be settled from the armchair. But there is an increasingly influential evolutionary argument that can be used to lend powerful support to EP.²²

¹⁹ My argument here, however, is compatible with views that ascribe value to some non-sentient entities so long as that value is very significantly lower than that of sentient beings (for relevant discussion, see Kahane (2014), 757–759). But on views that see ecosystems as the central locus of value in nature, and which regard wild animal suffering merely as a neutral aspect of the natural process, there isn't even space for the questions I'll consider here.

²⁰ Proto-mammals appeared around 200 million years ago. But fish, which are increasingly regarded as sentient, probably first appeared around 480 million years ago.

²¹ See Schopenhauer (1969).

²² I draw here heavily on Ng (1995); Horta (2015); Tomasik (2015) and Johannsen (2021). For criticism of this argument, see Palmer (2021).

This argument aims to show that the lives of most animals in the wild aren't worth living—that they are net negative. It starts from the background point that the evolutionary process inherently involves fierce competition for survival and reproduction. This is typically due to scarce resources, but even when resources become more abundant, evolutionary pressures dictate a corresponding expansion in population size all the way to the point of saturation—meaning that there is no escape from competition, with its inevitable losers—those who will suffer and die of starvation or, weakened by it, of disease or predation. This is bad enough. But there's the further point that such evolutionary pressures drive most species to the reproductive strategy of producing enormous numbers of offspring while investing little in any one of them (the 'r-strategy', which is used by many lizards, amphibians, fish, but is also common among small mammals).²³ The vast majority is expected to die, and die soon, without having a chance to reproduce themselves—in a stable environment, only about one of all those offspring will survive long enough to have offspring of its own (in many such species, the early-life death rate is greater than 90%²⁴).

Now the process of evolution involves large-scale death as part of its logic, and we saw Dawkins list the myriad of gruesome ways in which animals perish in nature. But while early death plays an important part in the argument, it doesn't directly bear on EP. The badness of death is usually understood as a function of the good it deprives one of; it makes (some) lives worse than they would have been. But arguably it doesn't add badness to the world.

What clearly does add badness to the world is suffering. Now pain evolved because of its adaptive benefits: because its motivational effects help reduce bodily harm, thereby increasing reproductive fitness. Much pain is thus instrumentally useful, bad in itself yet overall beneficial for the animal. But evolution doesn't 'care' who suffers or how much, and there's a great deal of suffering that is useless for the individuals who suffer it: pain is there because the capacity to experience it tends to promote reproductive success, not because it's useful in any given instance, and the disposition to feel pain is often triggered in contexts where it serves no function.²⁵ Moreover, evolution is entirely 'indifferent' to suffering that doesn't affect reproductive success. This is particularly true of the suffering associated with dying which doesn't affect reproductive fitness even when prolonged and intense. Evolution thus has no 'reason' to reduce that suffering—and death by predation, or by disease, can often take a long time.²⁶ If fish feel pain, then an extreme example is the salmon: most salmon's bodies collapse shortly after they have reproduced. And these are the most successful members of the species—the vast majority of salmon never reach that point.²⁷

So we have pervasive death, usually at a very young age, via starvation, dehydration, exposure, disease or predation, and the intense yet pointless suffering that such

²³ See also MacArthur & Wilson (1967); Pianka (1970); Jeschke et al., (2008).

²⁴ Vitt & Caldwell (2009), 138–139.

²⁵ See Dawkins, *ibid.*

²⁶ As Gary O'Brien pointed out to me, a mechanism that would selectively switch off such suffering would in fact be extremely costly, and selected against.

²⁷ See Horta (2010) for an attempt to calculate the astronomical amount of suffering associated with just one reproductive cycle of the Atlantic Cod in North America.

deaths typically involve. Animals who die in this way, just minutes, days or weeks after being born, have little chance to enjoy anything pleasant, certainly not anything likely to offset their painful death. Their lives are almost certainly not worth living.²⁸

If successful, this argument only directly establishes that a large majority of sentient lives in the wild have an overall negative level of well-being. While already a depressing result, this doesn't entail that the value of wildlife is net negative since the suffering that these lives contain is constrained by their brevity, and in principle that badness could be outweighed if the lives of the lucky surviving few were incredibly good. However, given the vast disparity in the numbers on either side, and the fact that even those who reach adulthood face numerous perils, this is unlikely; some argue that many of these 'luckier' lives might also be overall bad.²⁹ Moreover, that's so far only a claim about the overall value of routine life in the wild. Once we bring in the larger historical context, with its mass extinction events, and the massive suffering they involved,³⁰ and countless smaller catastrophes, as well as the more 'normal' processes via which entire species lose out and go extinct, there's far more badness in evolutionary history than suggested by the core argument. Finally, if we think about badness in aggregative terms, then given the staggering numbers in question, this badness could add up over millions of years to a vast deficit even if the local balance is just modestly negative.³¹

The strength of this argument for EP depends on where, in evolutionary history, sentience begins. If only mammals are sentient then the argument is somewhat weaker. But few deny sentience to birds and octopi,³² and the case for sentience in fish is robust.³³ Since fish exhibit the expansive reproductive strategy in its most dramatic form, the case for EP would be nearly conclusive if sentience is as widespread as that.

1.2 Climbing out of the red

Even if EP is true this cannot on its own show that the world is bad as a whole. If what we have in mind here is the entire universe, across space and time, then terrestrial evolutionary history, as vast as it seems to us, is a negligible portion of that unimaginably larger whole. However, if we accept the sentience condition, the cosmos is

²⁸ This, I believe, will almost certainly also be true if animal well-being included non-hedonic elements: an animal that dies shortly after coming to exist has little chance to exercise its species-typical capacities and could hardly be described as flourishing.

²⁹ Fischer (2022) argues not only that surviving r-strategists often lead lives not worth living, but that this is also true of many species who employ more conservative reproductive strategies.

³⁰ The Cretaceous-Paleogene extinction that killed the dinosaurs, a mere 66 million years ago, wiped out around 76% of all species; earlier extinction events were even more sweeping.

³¹ See Horta (2015). In fact, EP might be true *even if most animal lives are net positive*, since on critical-level views even net-positive lives add negative value to a population if they are below a certain threshold (see Broome (2004)).

³² See e.g. Low et al., (2012). Some even argue that invertebrates may be conscious; see Birch (2022) for discussion.

³³ See Brathwaite (2010).

largely an axiological desert, and the issue then turns on whether there is sentient life elsewhere in the universe, something we don't know.

I'll focus on the terrestrial case. Even here, EP obviously isn't the whole story. To begin with, the story isn't yet over and, for all we know, is just beginning. Now, if humanity will soon die out, and life will continue without any other intelligent beings replacing us, then the scope of EP would merely extend even further. The question, then, is what axiological difference is made by us humans. In large part this depends on our future and, since we don't know what it holds—whether our extinction is around the corner, or humanity (or, more likely, our post-human descendants) will go on for millions more years, and whether what follows is utopia or dystopia or something in between—we aren't remotely in a position to say whether pessimism holds even here on Earth.

Still, if EP is true then, when we humans first arrived on the scene, the planet was in a staggering axiological deficit, deep in the red. We can still ask what it would take to overturn this deficit, whether it's likely that human history so far has already managed to do so—or whether what we can call *Till-Now Terrestrial Pessimism* holds.

The idea that it's nature that is the problem, so to speak, and that it is we humans who might save the day, would be surprising to many; nowadays, I said, it's more common to see humanity as brutally usurping a harmonious natural order. Now if EP is correct, this rosy view of nature is deeply mistaken. This, however, is compatible with *Homo sapiens* making things even worse, and worse in awful new ways. Human history is a grim spectacle. It's natural to think that, although we are late arrivals, we have been extremely busy adding to the massive deficit we inherited.

However, this isn't clearly true. That human history is crammed with horrific episodes is indisputable. But again the question is whether these outweigh the good. Setting aside for now our unpleasant impact on other animals, it is plausible that the primary driver of value in human history is the value realized in individual human lives. If so, then our question becomes that of whether human lives were generally worth living. The considerations that support EP aren't at work here: the reproductive strategy of humans is to produce few offspring but invest massively in each. Still, human lives may nevertheless contain more bad than good. This is what is claimed by familiar pessimist views from Schopenhauer to Benatar.³⁴ On these views, the arrival of humanity did nothing to improve things; in fact, we would have added to the heap of awfulness *even* if we had been able to restrain our destructive tendencies. Till-Now Terrestrial Pessimism would be a given. But such pessimist views are widely rejected, and typically involve revisionist claims about human well-being; assessing such claims is beyond the scope of this paper. What I wish to highlight, however, is that if EP is true then Till-Now Terrestrial Pessimism is likely enough to also be true *even* if we hold on to far more flattering assumptions about humanity.

Parfit writes that when he considers the parts of human history he knows, it seems to him that they were overall positive, hedonically speaking.³⁵ Let's assume, for argument's sake, that this is broadly correct—that human history as a whole, while relentlessly grim, is nevertheless on the whole positive. This, however, hardly suffices to

³⁴ Schopenhauer (1976); Benatar (2017).

³⁵ Parfit (2011).

show that human history even begins to overturn the enormous deficit generated by many millions of years of animal agony. The issue here isn't so much that human history is just a brief late interlude in the larger terrestrial story. What matters are the numbers, and we humans are obviously massively outnumbered.

I assumed above that animal well-being is to be understood in classical hedonic terms. Hedonists hold that human well-being should also be understood in this way. If only simple hedonic states possess value then I think it's clear that, on EP, we won't even begin to overturn the deficit on even the rosier reading of the human past. Now there is admittedly still the future. But things will need to go incredibly well, for a long time, for us to finally climb out of the hole. On hedonism, the prospects of global pessimism are high.

Parfit, however, further holds that humans can also realize non-hedonic goods. Examples of such goods include things like deep personal relationships, virtue, and knowledge—and presumably also Williams's 'achievements and glories' of 'art, self-understanding, [and] nobility of character'.³⁶ These are often claimed to be not only objectively good, but also superior to sensory pleasure. This view is far more flattering to us humans. If we accept it, then the appearance of humans, and the emergence of civilization, mark a fundamental axiological shift in the history of the planet. A massive shift upwards, despite the accompanying depredations of constant warfare, tyranny, and enslavement. Call this view *Rational Exceptionalism*.

Rational Exceptionalism at least opens the door to the overturning of the deficit. It could most easily overturn it if we hold that such higher goods are lexically superior to simple hedonic states. However, this seems implausible. Even setting aside the familiar worries about lexicality,³⁷ it intuitively seems wrong to think that, say, as soon as the first humans placed a handprint on a cave wall, or begun to study the stars—or for that matter, Michelangelo painted the Sistine Chapel—hundreds of millions of years of suffering were overturned in a flash. And this isn't so surprising given that the intuitive plausibility of lexicality rests on comparisons between 'higher' and 'lower' goods. It seems less attractive when we compare higher goods and suffering; a screening of *La Dolce Vita* cannot make a never-ending life of unremitting suffering worth living—or even a pretty short one.

It is more plausible to think of Rational Exceptionalism as claiming that higher goods are just *much more* valuable than hedonic ones. Even if we accept this—though recall the limits suggested even here by the *La Dolce Vita* example—the amount of suffering that needs to be overturned is so staggering that it seems far from obvious that anything we humans have achieved so far—even ignoring our shortcomings and abject failures—would be sufficient to outweigh it.

There are a number of ways in which things might be even worse. There is first our gruesome impact on other animals. Admittedly until recently the number of animals affected was miniscule compared to the number who suffered throughout evolutionary history, and while factory farmed animals are both numerous and live horrific lives, their appearance seems to me far too recent to make a dramatic difference to the

³⁶ See also the discussion of what Sidgwick called 'ideal goods' in Parfit (1984), 454.

³⁷ See Arrhenius (2005).

balance so far. But how we go on treating domesticated animals—and perhaps even wild ones—could make a great difference to whether global pessimism ends up true.

And while things are pretty bad even if we think of the values of wholes in straightforward aggregative terms, it would be considerably harder to overturn the deficit on some other approaches. For example, Parfit suggests that we cannot simply balance goods and bads if these occur in different lives.³⁸ Parfit does hold that suffering that wasn't compensated within a life can still be outweighed—he gives the example of “people whose lives were worse than nothing, but whose lives were not very bad, since they did not involve long periods of intense suffering,” and suggests that if there were at least a hundred people whose lives were very good for each such life, that would outweigh that uncompensated badness. And Parfit speculates that human history was worth it even by this demanding standard—although, since human history contains many lives that were far more awful than the bad lives Parfit describes, it's unclear how many wonderful lives would be needed to outweigh them. In any event, if EP is correct, there were *quadrillions* of animal lives not worth living, vastly outnumbering all human lives. Unless it only matters whether suffering is uncompensated when it's experienced by persons, on such a view we are even further from overturning the deficit. Views that give greater weight to pain compared to equivalent pleasures,³⁹ and prioritarian views on which harms and benefits to those who are worse off matter more,⁴⁰ will yield similarly discouraging results.

If EP is true, then I think it is rather doubtful that human history has so far overturned this massive deficit. For this to even be on the table we need to have an incredibly rosy picture of the human past, to reject hedonism and axiologies that accentuate the negative, and to hold a strong form of Rational Exceptionalism. Conversely, however, even if we reject some or all of these assumptions, it remains the case that we aren't in a position to rule out that things will still come good in the long-term future, though this requires a heavy dose of old-fashioned optimism—hoping that our descendants will do incredibly well, and go on for a very long time, while being much nicer to non-human animals.

Over many millions of years, countless dazzling living things have evolved via natural selection. But evolutionary evil means that it's likely that, at least until our arrival on the scene, this just wasn't worth it. It is far from clear that humanity's emergence has improved things, and the future is unknown. There is thus a serious chance that global pessimism will end up true, that none of this will have been worth it in Parfit's sense. And if the world is overall bad, then it is worse than nothing. It would have been better if no sentient life had emerged, or even if there were absolutely nothing instead of (this) something.⁴¹

³⁸ Parfit (2019), 611.

³⁹ Hurka (2010); Benatar (2006).

⁴⁰ Parfit (1997).

⁴¹ If there's no sentient life anywhere outside the Earth then what happens here on little Earth decides the overall value of the cosmos (Kahane, 2014, 2021c). But it seems unlikely that we are alone (Westby & Conselice, 2020). If so, then our terrestrial drama makes only a negligible difference on this grander scale. Still, if evolution always follows the same brutal logic, then the amount of evolutionary evil out there is many magnitudes larger. And it seems plausible that if EP holds on Earth, it also holds generally. There is thus a serious chance that the universe is overall bad if EP holds.

2 Could it really have been better?

2.1 Necessary evils

It would be depressing if we concluded that the world as a whole is bad. But it's too early to get upset about that. We just don't know. However, we'd have little to cheer even if we knew that things *are* overall good. It is instructive that when we consider great evils, we rarely focus on questions about the value of wholes. The year 1944 obviously wasn't the greatest of times. But when we look back with horror at that time, this doesn't reflect a judgment about the overall value of the world. One reason we don't often form such overall evaluations is that it's just too hard to say. But another is that it often matters more whether things could have been better. Even if the world was overall net positive in 1944, this won't change the fact that something extraordinarily awful was happening within it, something awful that didn't have to happen. This is why, when we consider some evil, we want to know not just how it relates to larger units, but also whether it was avoidable, and what were the alternatives.⁴²

When horrors lie in the past, nothing we can do now will undo them. The remaining question for us is how to evaluate them, and what attitudes to have towards them.⁴³ Considered on their own, past horrors call for, well, our horror. They also call for our regret: they give us reason to wish the horrors had never happened, that things had been better—even if things were overall good. Actually, to speak about things being better is misleading since this doesn't distinguish between ways that good things might have been even better from ways awful things could have been avoided, or less bad. It's the latter that raises the question of whether things weren't worth it in Williams's Nietzschean sense—whether the past merits our regret even if the total good in it is greater than the bad.

Williams famously suggested that the 'constitutive thought' of regret is... "how much better if it had been otherwise".⁴⁴ But we cannot regret that things hadn't been better—or even less bad—in just any conceivable way. In most cases it is easy to conceive of superior alternatives to the actual. But if these are utterly fantastic, or even just far-fetched, the contrast with them cannot ground regret. They are simply too distant, they weren't really on the table. Contrast that with the instant regret we feel when we see the train leaving just as we arrive at the platform. Here we feel that things *really* could have been otherwise, that this better alternative was within reach—yet now lost. Regret contrasts the actual not with just any alternative, but only with realistic possibilities—only with what was a sufficiently likely (or near) counterfactual.⁴⁵

⁴² The alternatives matter even if we are confident that the world *is* overall bad since, although such a world is worse than nothing, there may still be numerous realistic alternatives that *would* have been overall good and therefore merit our attention more than the distant (and probably unrealistic) possibility of a complete void.

⁴³ Kahane (2021a).

⁴⁴ Williams (1981).

⁴⁵ See Wallace (2013), 72 ff.; Smilansky (2013), 3; Kahane (2019). See also Williams (2009a), 50.

But even when things could have realistically been better, there's often a complication. As Nietzsche reminds us, evil is often a causal condition for later good. And this means that when we consider those realistic alternatives, we often find that what would have been better in one respect, by not containing that evil, would also have been worse in another, by lacking the downstream good or anything like it. And if the good sufficiently outweighs the evil then we may conclude that the evil, while still regrettable, is no longer something that we regret, all things considered. We may still wish, of course, that we could have had the good without that evil. But if that wasn't realistically possible then we no longer have the sense that something was missed, that things took a wrong turn at some point. The badness is still there, but now that the broader causal connections, and the modal surroundings, are in view, our dismay is muffled by our appreciation of the evil's necessary role in bringing about that greater good—a greater good that otherwise wouldn't have been realized.

In other words, the evil will have been worth it in the sense suggested by Williams: we come to regard it as a *necessary evil*—not literally necessary, nor merely instrumentally necessary for something else, but something that is soberly accepted because it plays an essential part in a larger causal complex that we affirm because it also contains something of greater value.⁴⁶ As we saw earlier, this is a stronger condition for regarding the past as having been worth it despite the awfulness it contains. On Parfit's sense, to say things had been worth it is only to say that they are good on the whole. This, we saw, is compatible with feeling deep regret about the way things are. By contrast, to say that things were worth it in Williams's modal sense is to say, not merely that we prefer the way things actually are to nothing, but that we *also prefer them to all realistic alternatives from which the evil would have been absent*, because, by also lacking the dependent greater good, these alternatives would have been worse on the whole.⁴⁷

Williams writes approvingly of Nietzsche's refusal to forget that "suffering was necessary for everything that he and anyone else valued".⁴⁸ Without the slaves toiling in the Mines of Laurion, there would be no Parthenon, and no Socrates. Without rigid hierarchy and the surplus produced by brutal mass labor none of the great early civilizations would have emerged—and therefore, also none of what later followed.⁴⁹ But the causal chain doesn't stop there. When Dawkins exhorts us to be thankful for the evolutionary process, he is clearly gesturing at even more distant causal preconditions for our existence. If there weren't this evolutionary process—or even if it had

⁴⁶ As Williams (2009a) writes, "the thought that despite the horrors which underlie every human achievement... the enterprise will have been worthwhile... must be that the horrors were necessary... we can reflect 'without this, that could not be, and the value of that means that this, after all, was worthwhile.'" My proposed account of necessary necessary evil is meant to echo, but is different from, the conditions for regarding an evil as non-gratuitous in discussion of the problem of evil (see e.g. Kraay 2016). For related discussion, see Kahane (forthcoming).

⁴⁷ When evils are necessary in this way, all the (realistic) alternatives in which they don't occur are worse (or equally bad). But this needn't mean the world is the best of all realistically possible worlds: there may well be even better alternatives—alternatives that *still* contain that evil, plus *more* good.

⁴⁸ Williams (2001).

⁴⁹ See Scott (2017).

gone just a little differently—there would no *Homo sapiens*, and thus no Parthenon or Bach. But, as Dawkins is well aware, this process was far from benign.

It can seem indecent, I said, to regard evolutionary evil as necessary simply because it was required for mediocre us to arrive on the scene—even if evolutionary history is overall net positive, or if the value contributed by humanity—now or at some hoped for future point—outweighs all that past evil. Even if those conditions are met this still won't show that this evolutionary evil was necessary. What is at issue, then, is what exactly were the realistic alternatives to the way things had actually unfolded.

2.2 Could it have been better for past animals?

We can start by asking whether things could have been better for the multitude of past non-human animals that were in fact born, suffered, and died. Our focus here is on the way various counterfactuals would have benefited, or harmed, actual past animals—on so-called *person-affecting* differences in value.⁵⁰ But since the vast majority of animals aren't persons, I'll call this the individual-affecting perspective. As I'll understand it, such a perspective ranks possibilities in terms of the amount of aggregate well-being that the actual individuals who exist in them enjoy.⁵¹

Let's first consider the possibility that EP is false. Two things constrain realistic counterfactual improvements in the lives of actual past animals. First, animals compete for limited resources, and the interdependencies within and between species mean that improvements in the life of some animals will inevitably lead to even greater harm to others—when gazelles evade capture, lions starve. And we saw that the evolutionary pressure to reproduce means that thriving species will quickly go on to multiply up to the point where fierce competition for resources resumes, leading again to large-scale harm. Second, there's the familiar point that even small changes in the timing of procreation, let alone changes in which animals mate, will lead to different individuals coming to exist. Think, for example, of a herd of gazelles that faced a severe drought in actual history. In a counterfactual in which the weather is more forgiving, these gazelles will reproduce in different times, and with different partners, meaning that soon enough the individuals that comprise this herd in that clement alternative would all be different from the gazelles that had actually existed. Even if those possible gazelles flourish, this wouldn't have benefitted the actual gazelles that they replaced, meaning that this improvement won't count as better from an individual-affecting perspective.⁵²

In the animal context, such a 'non-identity effect' operates faster than in the human case, because most animals reproduce early, more often, and in much larger numbers. Any counterfactual change that is large-scale enough to count as a significant

⁵⁰ Parfit (1984).

⁵¹ There are other person-affecting approaches (see Greaves 2017). For example, *presentist* approaches consider only presently existing individuals while *necessitarian* approaches consider only individuals that exist in all the possibilities in question. But concern for evolutionary evil makes no sense on the presentist approach while, for reasons that will emerge below, the necessitarian approach rules out virtually all comparisons between different evolutionary counterfactuals.

⁵² This is an instance of what Parfit called the 'non-identity problem' (Parfit, 1984), as applied to historical events (see Kahane 2019).

improvement in the lives of enough past animals would also be one from which all later actual animals are ‘erased’. While these ‘absent’ animals no longer suffer as they did in actual evolutionary history, neither do they go on to live lives not afflicted by suffering. These later animals obviously don’t benefit from the change and, at least on some views, they lose out.⁵³ In addition, for it to be even possible to benefit a given temporal slice of past animals, we need to hold fixed all prior suffering since otherwise the ‘benefited’ animals won’t even come to exist.⁵⁴

So only a sliver of past animals could have enjoyed better lives. As we saw, such benefits would typically come at a cost, often great cost, to other contemporaneous animals, so it’s not even obvious that animals existing at that point in time would, in aggregate, benefit overall. However, while the logic of the evolutionary process leaves little room for significant improvement, even fairly small changes in certain extrinsic factors *could* make a great difference. Think, in particular, of mass extinction events like the one that wiped out the dinosaurs. Perhaps it could have been avoided had an asteroid taken a different course. *If* animal lives are on the whole positive, as we’re now assuming, then the removal of an extinction event would have led to a dramatic improvement in the lives of all the animals that perished in that catastrophe—and utterly change the course of evolution.

But we also need to ask what, from such an individual-affecting perspective, would count as the *overall best* alternative to the actual world. This would depend, first, on whether we see the erasure of everyone following a point of ‘improvement’ as worse for those erased. Even if we reject this idea, surely alternative courses of evolution from which more actual individuals are ‘erased’ contain less individual-affecting value. And this already pushes us towards favoring more recent evolutionarily changes—meaning that all prior evolutionary evil remains in place.⁵⁵ To complicate things further, there’s also us humans, past, present and perhaps future.⁵⁶ And there were also ways in which great suffering to past humans were avoidable.⁵⁷ When we combine Rational Exceptionalism with the earlier point about favoring later changes, it may be that the best individual-affecting alternatives would be ones where the point of improvement occurs within *human* history.

Things are different, though, if EP is true. In terms of ‘offering’ past animals lives that are better, the same constraints hold, with the added worry that we are merely making horrible lives just a bit less awful. This would be clearest in the case of mass extinction events since the ‘normal’ that such disasters prevented was precisely the more common nastiness of numerous lives not worth living. In fact, if EP is true, I’m not so sure that mass extinction events really made most of their victims worse off.

⁵³ See e.g. Holtug (2010); Arrhenius & Rabinowitz (2015).

⁵⁴ I discuss this point as applied to past humans in Kahane (2019).

⁵⁵ See Kahane (2019). On another way of developing the individual-affecting view, we should prioritize earlier individuals over later ones. But this would imply that we should favor even slight improvements in the lives of the first sentient beings at the expense of all the rest. This seems implausible.

⁵⁶ It’s unclear whether and how future individuals figure on such a perspective. Arguably, it should encompass whichever future people and animals would in fact go on to exist, even if, right now, that’s not a fixed set of individuals.

⁵⁷ Smilansky (2013); Kahane (2019).

But EP would *also* mean that when we favor possible improvements to the animals who had existed at a certain time slice, the consequent erasure of all the quadrillions of animals who followed is no longer neutral, or even a disadvantage, but a clear benefit to the majority of those erased animals—animals who would be better off not existing. So while it's still the case that only a sliver could have led better *lives*, it would also be true, on EP, that things nevertheless *could* have been better for the vast majority of past animals.

Given these points, and since it's hard to see how it could be overall better to hold an immense number of lives not worth living fixed in order to secure a modest later benefit to a single generation, it seems to me that EP would likely mean that the best individual-affecting scenario for past animals isn't one where some animals go on to lead lives less afflicted by suffering, but a scenario where things take a different turn early on, meaning that *none* of the sentient beings who had actually existed remains.

In this grim way, things *could* have been significantly better for past animals on a much grander scale *only* if EP is correct. There are, however, two ways in which all actual past animals could have been erased. On the first, evolution takes a somewhat different route before the first sentient being appears. But then we're just replacing these quadrillions of miserable lives with *other* quadrillions of miserable lives. We have just 'shifted' that past evil from actual to possible animals.⁵⁸ Even if we give special weight to individual-affecting considerations, it seems wrong to just ignore the misery of those alternative animals, or that this could address our horror about evolutionary evil. There is, however, a way to avoid this unpleasant upshot: an alternative history in which *no* sentient life ever evolves. This, I believe, is what EP would point to, if we focus exclusively on the interests of actual animals. Whether adding us humans to the picture can change the overall verdict is a question I'll return to below.

2.3 Could it have been significantly better, impersonally speaking?

We just saw that it's implausible to give weight only to whether some counterfactual change benefits or harms actual past individuals. We should also be concerned with 'impersonal' value: with how much pain (and pleasure) a given course of evolutionary history contains, regardless of which individuals it contains. There are, of course, also those who anyway think that what really matters is how much suffering there is, not who experiences it,⁵⁹ and some hold that how specific individuals fare matters only if they are persons.⁶⁰

So consider next whether there is greater space for things to have been better, with respect to animal suffering, if we look at the past in this impersonal way, where possibilities are ranked by the total amount of value their contain.⁶¹ We again start

⁵⁸ Kahane (2019).

⁵⁹ See again Parfit (1984).

⁶⁰ Such a view is developed by Jeff McMahan (personal communication) in unpublished work, and suggested by the common talk of *person*-affecting morality, not to mention person-affecting views that directly appeal to personhood. For discussion, see Purves & Hale (2016).

⁶¹ Kahane (2019). Impersonal views that focus on average value are notoriously unable to give weight to numbers. To simplify the discussion, I'll also set aside prioritarian or pain-prioritising impersonal axiolo-

by considering the possibility that EP is false. It seems plausible that evolution could have taken quite different paths and that, of these, at least some would have contained less suffering, and overall more good, than actual evolutionary history. Since we are no longer concerned with which individuals exist in the alternate course of evolution we are considering, we can ignore the non-identity effect and are therefore no longer constrained to improvements at a given time—we can now add up such improvements, allowing for much larger scale improvements.⁶²

How much better some of these courses of evolution could have been is hard to say. Perhaps the range of options is fairly rigid and things would inevitably take a fairly similar shape,⁶³ meaning that even on an impersonal approach the range of possible improvements would still be limited. And even if the range of options is wider than that, it's at least possible that actual evolutionary history is fairly close to the optimum realistically possible. However, to the extent that at least some mass extinctions needn't have occurred, it seems plausible that there was space for significant improvement.

But even if there were such counterfactuals involving significantly less animal suffering, we cannot assume that an overall impersonal calculus would favor such alternatives. After all, such a standpoint is concerned only with the amount of value in different possibilities, not with who realizes that value. But if so, it's hard to see why our tally must take into account only the well-being of animals. It should also include the value that might be added later down the line by possible rational beings—and even more so on Rational Exceptionalism.

Now, this broadening of scope isn't likely to favor the course of evolution that led to us humans. Admittedly, many of the things we dislike about humanity have their source in evolutionary pressures that would have similarly shaped alternate intelligent lifeforms. It's even possible that the only realistic route to intelligent life goes through something like the apes and looks a lot like us. Still, I suspect that we humans aren't close to being the best that could have evolved. And if different intelligent lifeforms could have evolved that would have been nicer and smarter than us—or simply ones that would have had a brighter long-term future—then it's the value potentially contributed by such alternative non-human civilizations that's likely to count most in the impersonal calculus. By contrast, an alternate evolutionary path in which there is considerably less suffering, but which *doesn't* lead to any rational lifeform is likely to rank rather low (imagine the dinosaurs getting spared, and mammals therefore never evolving). And I don't think we can rule out that an evolutionary history that is overall *even bloodier* than ours yet leads to an especially successful rational lifeform would be favored (for example, harsher conditions and increased competition may increase selection pressure for higher cognitive capacities and ingroup cooperation).

So far as I can see, we get the same broad result even if EP is true. From an impersonal standpoint, what ultimately matters is the overall value of an entire course of history and, as we saw, EP cannot decide this overall value in scenarios in which

gies of the sort mentioned earlier. I think my overall argument will also go through on such views, although they set a higher bar for overall positive outcomes, as well as make them less probable.

⁶² See Kahane (2019).

⁶³ Conway Morris (2003).

intelligent lifeforms emerge. If Rational Exceptionalism is true, or, even if it's false, if an intelligent civilization is sufficiently prosperous and continues long enough, then EP can be ultimately outweighed. If that's the case, then from an impersonal standpoint we shouldn't hesitate to regard even EP as a necessary evil.

2.4 Evolutionary evil as necessary evil

We started out asking how evolutionary evil bears on the overall value of larger units. We have now also mapped some of the 'sideways' axiological terrain: ways in which things could have been realistically better, or less bad, on various scenarios. We saw that the very nature of the evolutionary process places severe limits on how better things could have been, so far as animals are concerned, and that if we care about the actual past animals who had suffered, there was only limited space for things to have been better for them, beyond the erasure of overall bad lives. I now turn to ask how, in light of all that, we should regard evolution, and the suffering associated with it.

We can begin by asking how we should regard evolution when we adopt a purely impartial perspective, whether individual-affecting or impersonal, or, as I favor, some combination of the two.⁶⁴

Some realistic alternatives to actual history may have contained less animal suffering. But we saw that there is no realistic scenario on which the evolutionary process rolls on, and sentient, and perhaps also intelligent, beings emerge, yet where there was no massive animal suffering—suffering on a grand scale is simply inherent in the logic of evolution. So the first choice we need to make within a naturalist framework is between a world that contains a great deal of evolutionary evil, and one in which the evolutionary process never gets going, or at least somehow never leads to sentience.

Let consider first this 'zero' option. It's possible that it was extremely fortuitous that life emerged on Earth. Whether it was a genuine possibility that no sentient life whatsoever would ever appear in the entire universe is harder to say. On some views, the emergence of sentient life is a virtual inevitability.⁶⁵ If so, then even if one grimly concludes that it would be preferable if there weren't any sentient life at all, this just won't be feasible. Evolutionary suffering would literally be a (physical) necessity. But assume that a universe without sentience really was on the table. Could this really have been for the best?

We're assuming that sentience is necessary for value. Since in the naturalist framework there is no path to sentience without the evolutionary process,⁶⁶ this would be world containing nothing that matters.

From an impersonal standpoint, we should prefer nothing—or such a barren world—to the actual world only if global pessimism were true. But even if global pessimism were true this *won't* entail that, from such an impersonal perspective, this

⁶⁴ Kahane (2019).

⁶⁵ Conway Morris (2003); Davies (2006). See also England (2013), though England argues for the near inevitability of life, not specifically of sentient life.

⁶⁶ A possible exception are so-called Boltzmann brains—brains momentarily assembled out of random fluctuations of matter in a vast enough universe.

would be the possibility we should most favor. Global pessimism only means that the actual world is worse than nothing, not that nothing is the best that could (realistically) be. After all, there might still be alternative realistic counterfactuals in which immense amounts of value is realized—scenarios in which evolution takes a different course and alternative intelligent beings go on to do great things. From an impersonal standpoint, we should wish that no sentient life ever emerged only if global pessimism applies to *all* realistic scenarios. I said earlier that we're not in a position to say whether global pessimism will be true of the actual world; we're *very* far from being able to assert such a modally ambitious form of pessimism.⁶⁷ This is so *even* if we're confident that EP is true not just of the actual world but of all evolutionary scenarios.

From an individual-affecting perspective, there will also be no reason to favor the zero option if EP is false. If most animal lives were worth living, such a scenario offers them nothing, and perhaps amounts to a loss. It's more likely, though, that EP is true. We saw that this would mean that it *would* have been better for most past sentient beings not to have existed and, if we don't want to replace them with other miserable beings, this could only be realized via this 'zero option'. Within a purely individual-affecting perspective, the question becomes that of whether the relief-by-erasure of all those quadrillions of past animals outweighs the loss of rational beings and the distinctive higher goods associated with them.

While we're now considering how things could have been, this question is nevertheless parallel to the question we asked earlier, when we considered whether what humans bring to the table can erase the axiological deficit. I suggested there that the likelihood that the value humans contribute will outweigh the prior darkness is much higher if we accept Rational Exceptionalism, and hold on to the prospect of a long and rosy future. I concede that incorporating the future in an individual-affecting view isn't straightforward, but if we also give weight to impersonal considerations (as I believe we should), then we *must* also take into account that possible bright future, as well as even rosier counterfactuals involving other intelligent beings.

I don't rule out that when we fully think through the significance of Darwinism, we must conclude that it would have been better if evolution never got going in the first place. But it seems to me that when what is at stake is there being nothing of value whatsoever, and given that there's simply no alternate route to realizing anything good in our kind of world, and so long as we cannot rule out that the rational beings (whether us, or like us, or far better) that follow (or could have followed) the millions of years of pre-rational agony might turn things around—then we should at least provisionally reject the zero option even when we consider things from a grand impartial perspective.⁶⁸

⁶⁷ One familiar solution to Fermi's paradox suggests a way in which such pessimism could be true: if all sophisticated civilizations quickly destroy themselves, and their planets, as soon as they develop nuclear capacities.

⁶⁸ Would we acquiesce, in this way, in a thousand-year Reich in which billions of humans were brutally murdered if this were necessary for a glorious future containing greater good? If not, how can we decently accept millions of years of *quadrillions* of suffering animals? I admit I struggle with this question. I think it helps that the vast majority of evolutionary evil is contained in brief lives, limiting the amount of evil each contains—arguably an instance of the 'reverse repugnant conclusion' (see Holtug, 2022). So a better human analogue would involve infant mortality. Yet although there were high infant mortality rates

Now if we reject the zero option, or think it wasn't a realistic possibility, there will still be many possible courses of evolutionary (and broader) history on the table to choose from. But *all* of these will contain quadrillions of suffering animals. So on both the impersonal and individual-affecting standpoints, that suffering will be seen as a necessary evil, as a necessary route to a greater good.

This hardly means we should regard evolution with unqualified gladness. But while retaining a sense of horror, we also shouldn't regard evolution with regret. Still, so long as we view things from an impartial standpoint, we aren't resigning ourselves to evolutionary evil because it improbably led to *us* humans, but because it's a causal precondition for *anything* good, including for rational beings *like* us—but which, in the impartially favored scenario, where the good outweighs the bad, almost certainly won't *be* us if we give impersonal considerations significant weight. We therefore needn't resign ourselves, from an impartial perspective, to evolution unfolding in the specific way it actually did. We don't regret that evolution took place in some form. But we should probably regret that it unfolded in *this* specific form. We don't wish away evolutionary evil, but we likely wish away *Homo sapiens*.

But our own perspective isn't this impartial. So I now finally turn to ask whether, from our collective human perspective, we *are* permitted to endorse this specific path leading to us. That of course is what we would prefer if we regarded things from a purely self- (or rather species-) interested standpoint. But I'm asking whether such an attitude could also be morally defensible—defensible *in light* of the verdict of a purely impartial outlook. In particular, the question is whether we are permitted to prefer a suboptimal actuality that includes us humans to impartially superior alternatives that don't—a kind of deontological prerogative allowing us to prefer, from our partial perspective, something that is less than the (impartially) best.

If we endorse the zero option from an impartial standpoint, then I doubt we can then insist, when we adopt our much narrower human perspective, that we still prefer the awfulness as it was simply because it led to us; but we're now setting aside the zero option. Even if we reject the zero option, we couldn't decently prefer the actual course of things if we think that global pessimism is true of the actual world; but, again, we're not in a position to confidently assert that. So the question is no longer *whether* to endorse evolutionary suffering as a necessary evil but only what *form* it should take. Can we decently prefer actuality to alternatives which contain less animal suffering?

But consider the following. Starting first with the individual-affecting standpoint, we saw that there are severe constraints on how better things could have been to actual past animals, and, if EP is false, they could have been better only for a sliver.⁶⁹ Even if that improvement was overall favored from an individual-affecting perspective, such an improvement—while certainly not to be sniffed at—would no longer involve the overwhelming numbers that would swamp whatever value humans bring to the

throughout human history, we don't think that this, on its own, means that human history wasn't worth it, even if we suppose that many of those brief human lives were net negative (though admittedly the number of humans who died this way is comparatively tiny).

⁶⁹ While if EP is true those awful lives would just be replaced by other awful lives if we reject the zero option.

table. So it's not obvious that such an improvement will be enough for us to wish humanity away. Moreover, we also saw that an individual-affecting standpoint may well overall favor later improvements that occur in human history—meaning that all prior evolutionary suffering is held fixed. If that's the case, then the impartial and human standpoints could coincide.

Consider next the impersonal perspective—and I suggested that we must give at least some weight to impersonal differences in value even if we give outsize weight to actual individuals.⁷⁰ While we saw that here there's more space for impersonal reduction in the suffering animals would endure over the course of evolution, we also saw that we cannot be confident that the impersonally best alternatives would really involve such a reduction; and for all we know these alternatives might contain *more* animal suffering. Rather, the key factor in making some alternative impersonally superior to actuality is likely our replacement by more impressive intelligent species.

I don't think we are required to prefer such a superior counterfactual. Perhaps we 'owe' something to the numerous suffering animals to which we owe our existence; but we owe nothing to those hypothetical intelligent beings.⁷¹ And while it's hard to reject a *less awful* alternative just because we'd be erased from it, I don't think we're at all required to prefer an alternative simply because it would be overall better—especially if the key component in making the alternative superior is our replacement (must we regret that our parents didn't have children better than us?) and when, for all we know, that alternative might be worse in terms of animal suffering.⁷²

So I think that there's a provisional case for regarding evolutionary evil in general, and perhaps even the specific path evolutionary history in fact took, as a necessary evil. We cannot regard it with the enthusiasm that Dawkins's remark suggests. But neither must we overall regret that such suffering happened, and perhaps even that it happened *this* way.

3 Conclusion

To discover that we weren't created in seven days but are the products of a blind evolutionary process over millions of years is also to discover that the past contains an extraordinary amount of suffering. Looking back, we can ask whether those millions of years of agony were really worth it. One way in which that suffering wouldn't have been worth it is if the balance of value in the world is negative. I've argued that there's a strong (if inconclusive) case for thinking that terrestrial evolutionary history is so awful that it's on the whole bad, and that this would generate an enormous axiological deficit that would be difficult to overturn even if we cling to some of the most optimistic assumptions about humanity.

⁷⁰ For discussion of views that combine individual-affecting and impersonal considerations, see Ross (2015).

⁷¹ Williams (2009c) similarly defends prioritizing humanity over (actual) superior intelligent beings.

⁷² I concede that if we could be confident that this counterfactual would contain far less animal suffering than the case for holding on to the actual would be weak.

Even if we, or our better future descendants, will overturn this deficit, those howls and shrieks, going on for hundreds of millions of years, remain a blot on the past. Yet without this awfulness we would never have come to exist. It would be monstrous to look back with pleasure on the evolutionary process simply because all that evil was a precondition for our existence. But evolutionary evil is a causal precondition, not just for our own existence, but for the realization of positive value of any kind in our kind of universe: there's either millions of years of agony or nothing of value at all. And while I admit I remain uncomfortable about this conclusion, I have argued that, at least so long as there's a decent chance that the overall balance would be positive, we can nevertheless justifiably regard the grim evolutionary past as a necessary evil and, therefore, as having been worth it. We must, however, take seriously the crushing possibility that the actual past, and even all realistic evolutionary alternatives to it, are worse than nothing. But so long as we don't endorse this bleak view, we must conclude that, while our universe is obviously far from being the best of all possible worlds, evolutionary suffering is an inevitable aspect of the best of all *realistically* possible worlds.

This is admittedly itself a depressing fact about the cosmos. But we can only sensibly regret realistic alternatives to the way things are, and, from a naturalist standpoint, there's no realistic alternative to the naturalist framework itself. We can still fantasize about more distant counterfactuals, we might still long, perhaps, for the cuddly comforts of theism.⁷³ But even that wouldn't have offered real relief. Now, that imagined universe would have probably been far rosier. A loving God wouldn't have permitted all that suffering. And while creationism is preposterous as an account of the actual world, it's at the same time also axiologically vastly superior to Darwinism. Still, why even assume that, had God existed, He would have created narwhals and aardvarks in the first place? And had He existed, God almost certainly wouldn't have created any of the individual animals who had actually existed—and it's also incredibly unlikely that He would have created us humans, as opposed to the numerous superior alternatives. Thus, so far as animals and suffering are concerned—and so far as *we* are concerned—that distant rosy counterfactual is as good as nothing.^{74,75}

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

⁷³ Though see Kahane (2011).

⁷⁴ For further spelling out of this last argument see Kahane (2022) and Kahane (forthcoming).

⁷⁵ I am grateful to two anonymous referees for very helpful comments, and especially to Gary O'Brien for his excellent and extensive comments on several drafts. I also benefited from comments from participants in a workshop in Reading in July 2019. While I don't accept the intriguing argument in Nagasawa (2018), thinking about it has also helped me sharpen my argument here.

References

- Arrhenius, G. (2005). 'Superiority in Value'. *Philosophical Studies*, 123(1–2), 97–114
- Arrhenius, G., & Rabinowicz, W. (2015). 'The Value of Existence'. In I. Hirose, & J. Olson (Eds.), *The Oxford Handbook of Value Theory* (pp. 425–443). Oxford: Oxford University Press
- Benatar, D. (2006). *Better Not To Have Been Born*. Oxford: Oxford University Press
- Benatar, D. (2017). *The Human Predicament*. Oxford: Oxford University Press
- Birch, J. (2022). 'The Search for Inveterate Consciousness'. *Nous*, 56, 133–153
- Brathwaite, V. (2010). *Do Fish Feel Pain?*. Oxford: Oxford University Press
- Broome, J. (2004). *Weighing Lives*. Oxford: Oxford University Press
- Conway Morris, S. (2003). *Life's Solution: Inevitable Humans in a Lonely Universe*. Cambridge: Cambridge University Press
- Davies, P. (2006). *The Goldilocks Enigma*. London: Penguin
- Darwin, C. (1859). *On the Origins of the Species*. London: John Murray
- Dawkins, R. (1995). *River Out of Eden*. NY: Basic Books
- Dawkins, R. (2010). 'The Greatest Show on Earth.' Paper delivered at the University of Auckland, New Zealand, 13 March. <https://www.alumni.auckland.ac.nz/en/photo-galleries-and-video/public-lectures/richarddawkins-thegreatestshowonearthlive/richarddawkins-thegreatestshowonearthlive-1.html>
- Dennett, D. (1995). *Darwin's Dangerous Idea*. NY: Simon & Schuster
- England, J. (2013). 'Statistical Physics of Self-Replication'. *Journal of Chemical Physics*, 139, 121923
- Fischer. (2022). 'Two Challenges to Johannsen on Habitat Destruction'. *Philosophia*, 50, 865–873
- Greaves, H. (2017). 'Population Axiology'. *Philosophy Compass*, 12(11), e12442
- Holtug, N. (2010). *Persons, Interests, and Justice*. Oxford: Oxford University Press
- Holtug, N. (2022). 'Prioritarianism and Population Ethics', Arrhenius, G., Bykvist, K., Campbell, T., Finneron-Burns, E. *The Oxford Handbook of Population Ethics*, 38–62
- Horta, O. (2010). 'Debunking the Idyllic View of Natural Processes'. *Telos*, 17, 73–90
- Horta, O. (2015). 'The Problem of Evil in Nature: Evolutionary Bases of the Prevalence of Disvalue'. *Relations Beyond Anthropocentrism*, 3(1), 17–32
- Hurka, T. (2010). 'Asymmetries in Value'. *Nous*, 44(2), 199–223
- James, W. (1929). *The Varieties of Religious Experience*. Cambridge: Cambridge University Press
- Jeschke, J. M., Gabriel, W., & Kokko, H. (2008). 'r-Strategists/K-strategists'. *Encyclopedia of Ecology*, 4, 3113–3122
- Johannsen, K. (2021). *Wild Animal Ethics: The Moral and Political Problem of Wild Animal Suffering*. London: Routledge
- Kahane, G. (2011). 'Should We Want God to Exist?'. *Philosophy and Phenomenological Research*, 82, 3: 774–696
- Kahane, G. (2014). 'Our Cosmic Insignificance', *Noûs*, 48, 4: 745–772
- Kahane, G. (2019). 'History and Persons'. *Philosophy and Phenomenological Research*, 99, 1: 162–187
- Kahane, G. (2021a). 'The Significance of the Past'. *The Journal of the American Philosophical Association*, 7, 4: 582–600
- Kahane, G. (2021b). 'Optimism Without Theism? Nagasawa on Atheism, Evolution, and Evil', forthcoming in *Religious Studies*
- Kahane, G. (2021c). 'Importance, Value, and Causal Impact', forthcoming in the *Journal of Moral Philosophy*
- Kahane, G. (2022). 'Is the Universe Indifferent? Should We Care?'. *Philosophy and Phenomenological Research*, 104, 676–695
- Kahane, G. (forthcoming) (Ed.). 'Should Atheists Wish That There Were No Gratuitous Evil?', *Faith and Philosophy*
- Kershenbaum, A. (2021). *The Zoologist's Guide to the Galaxy*. London: Penguin
- Kraay, K. (2016). 'God and Gratuitous Evil (Part I)'. *Philosophy Compass*, 11, 905–912
- Low, P., et al. (2012). *Cambridge Declaration on Consciousness*. <http://fcmconference.org/img/CambridgeDeclarationOnConsciousness.pdf>
- MacArthur, R., & Wilson, E. O. (1967). *The Theory of Island Biogeography*. Princeton, NJ: Princeton University Press
- Murray, M. (2008). *Nature Red in Tooth and Claw: Theism and the Problem of Animals Suffering*. Oxford: Oxford University Press

- Nagasawa, Y. (2018). 'The Problem of Evil for Atheists'. In N. N. Trakakis (Ed.), *The Problem of Evil: Eight Views in Dialogue*. Oxford: Oxford University Press
- Nietzsche, F. (2003). In Bittner (Ed.), *Writings from the Late Notebooks*. R., Cambridge: Cambridge University Press
- Ng, Y. K. (1995). 'Towards Welfare Biology: Evolutionary Economics of Animal Consciousness and Suffering'. *Biology and Philosophy*, 10(3), 255–285
- Palmer, C. (2021). 'The Value of Wild Nature', *Philosophia*. 50 (3):853–863
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press
- Parfit, D. (1997). 'Equality and Priority', *Ratio*, 10 (3): 202–221
- Parfit, D. (2011). *On What Matters* (II vol.). Oxford: Oxford University Press
- Pianka, E. (1970). 'On 'r' and 'K-' selection'. *American Naturalist*, 940, 592–597
- Purves, D., & Hale, B. (2016). 'Non-Identity for Non-Humans'. *Ethical Theory and Moral Practice*, 19, 1165–1185
- Ross, J. (2015). 'Rethinking the Person-Affecting Principle'. *Journal of Moral Philosophy*, 12(4), 428–461
- Schopenhauer, A. (1969). *The World as Will and Representation*. Vol. II, NY: Courier Dover
- Schopenhauer, A. (1976). *Essays and Aphorisms*. London: Penguin
- Scott, J. C. (2017). *Against the Grain*. New Haven, CT: Yale University Press
- Smilansky, S. (2013). 'Morally, Should We Prefer Never To Have Existed?'. *Australasian Journal of Philosophy*, 91(4), 655–666
- Tomasik, B. (2015). 'The Importance of Wild-Animal Suffering'. *Relations Beyond Anthropocentrism*, 3(2), 133–152
- Tomasik, B. (2019). 'How Many Wild Animals Are There?' in his online *Essays on Reducing Suffering*. <https://reducing-suffering.org/how-many-wild-animals-are-there/>
- van Inwagen, P. (2008). *The Problem of Evil*. Oxford: Oxford University Press
- Vitt, L. J., & Caldwell, J. P. (2009). *Herpetology: An Introductory Biology of Amphibians and Reptiles*, 3rd ed., 138–139
- Wallace, R. J. (2013). *The View from Here: On Affirmation, Attachment, and the Limits of Regret*. Oxford: Oxford University Press
- Westby, T., & Conselice, C. (2020). 'The Astrobiological Copernican Weak and Strong Limits for Intelligent Life'. *The Astrophysical Journal*, 896(1), 58
- Williams, B. (1981). 'Moral Luck' reprinted in his *Moral Luck* (pp. 20–39). Cambridge: Cambridge University Press
- Williams, B. (2001). 'Introduction' in Nietzsche, F. *The Gay Science*. Cambridge: Cambridge University Press
- Williams, B. (2009a). 'The Women of Trachis: Fictions, Pessimism, Ethics'. In his *The Sense of the Past: Essays in the History of Philosophy*. Princeton University Press. 49–59
- Williams, B. (2009b). 'Unbearable Suffering'. In his *The Sense of the Past: Essays in the History of Philosophy*. Princeton, NJ: Princeton University Press, 331–340
- Williams, B. (2009c). 'The Human Prejudice'. In J. A. Schaler (Ed.), *Peter Singer Under Fire: The Moral Iconoclast Faces His Critics* (pp. 77–97). Chicago: Open Court Publishing