

SPECIES' WELLBEING, COUNTERFACTUAL COMPARATIVE HARM, AND THE NON-IDENTITY PROBLEM

J. Spencer Atkins
Binghamton University
jatkings4@binghamton.edu

Abstract

The non-identity problem raises problems for many versions of the counterfactual comparative account of harm. If an individual's existence depends on climate change, then we cannot say that climate change makes this individual worse off than they would be otherwise, since otherwise they would not exist. However, I argue that consideration for species' wellbeing avoids the non-identity problem: the species can be worse off than it would have been otherwise because the species existence does not depend on climate change. I first examine several views of counterfactual comparative harm and argue that they are subject to the non-identity problem. Then I survey a number of views of species, showing that they are consistent with my argument. I, then, offer a novel account species' wellbeing and species' harm. Species harm and wellbeing is the aggregate projected aggregate welfare of all the individual members over time. I then argue that this account of species' wellbeing avoids the non-identity problem. In the last section, I answer objections.

Keywords

Non-Identity Problem; Counterfactual Comparative Account of Harm; Species' Wellbeing; Biotic Welfare; Species

Resumen

El problema de la no identidad plantea problemas a muchas versiones del análisis comparativo contrafáctico del daño. Si la existencia de un individuo depende del cambio climático, entonces no podemos decir que el cambio climático hace que este individuo esté peor de lo que estaría en caso contrario, ya que de lo contrario no existiría. Sin embargo, en este artículo sostengo que la consideración del bienestar de las especies evita el problema de la no identidad: las especies pueden estar peor de lo que estarían si no existieran porque su existencia no depende del cambio climático. En primer lugar, examino una serie de visiones del daño comparativo contrafactual y argumento que están sujetas al problema de la no identidad. A continuación, examino una serie de puntos de vista sobre las especies, mostrando que son coherentes con mi argumento. A continuación, ofrezco una nueva explicación del bienestar y el daño de las especies. El bienestar y el daño de las especies es el bienestar agregado proyectado de todos los miembros individuales a lo largo del tiempo. A continuación, argumento que esta explicación del bienestar de las especies evita el problema de la no identidad. En la última sección, respondo a las objeciones.

Palabras clave

Problema de la no identidad; Reporte comparativo contrafáctico del daño; Bienestar de las especies; Bienestar biótico; Especies

Introduction

Duncan Purves and Benjamin Hale (2016) argue that nonhuman organisms are subject to the non-identity problem. The non-identity problem proposes that many nonhuman organisms' existence is dependent upon climate change, a source of harm. Thus, they cannot really be said to have been harmed, since their existence is contingent on the allegedly harmful phenomena. Their argument threatens many versions of the counterfactual comparative view of harm (hereafter, counterfactual comparative).¹ The counterfactual comparative account states, broadly, that S is harmed if and only if S is made worse off than S would have been otherwise, e.g., without the harmful event. I argue that consideration for species harm avoids the non-identity problem. The existence of a species is not contingent upon the policies that, for example, caused climate change. A species might exist without being dependent upon climate changing policies, even if some members of the species are. Even if we cannot say climate change harms individual members of a species, we could say that the species overall is harmed. A single species, regardless of which individual members exist, might fare better or fare worse. We can, therefore, use the counterfactual comparative view to assess species harm ensuing from climate change. The species is not subject to the non-identity problem.

¹ By no means are they the only ones who argue against the counterfactual comparative account of harm. McMahan (2013) argues that the counterfactual comparative account cannot be a full account of a harm because it lacks explanation for non-comparative or intrinsic harms. Bradley (2012) argues that the counterfactual comparative account cannot make sense of omissions and failures to benefit. These objections, though interesting, are not the focus of this paper; I will focus exclusively on Purves and Hale's argument for Non-Identity for Nonhumans.

In the second section, I lay out Purves and Hale's argument against counterfactual comparative and briefly explain the non-identity problem. I then offer several versions of counterfactual comparative, showing that each of them is subject to Purves and Hale's argument. The third section explains the difficulty in defining *species*. After offering several views about the nature of species, I argue that each of them is consistent with my way around the non-identity problem. Next, I argue that species' wellbeing, i.e., species harm and benefit, is best understood as aggregate welfare of the individual members of the species plus the projection of wellbeing in the near to medium future. This plausible view of species' wellbeing, I argue in section five, gives us the tools to avoid non-identity for nonhumans. Purves and Hale's threat to counterfactual comparative fails. In section six, I address objections to my argument: (1) my argument cannot make sense of species who exist in virtue of climate change—the non-identity problem is still a problem—and (2) there is another version of counterfactual comparative that sidesteps the non-identity problem, so my argument is not necessary. I then conclude in section seven.

Purves and Hale and the Counterfactual Comparative Threat of Non-Identity

Many philosophers have argued that the counterfactual comparative account of harm is likely the most plausible view of harm available.² Duncan Purves and Benjamin Hale, however, challenge this conviction for nonhuman animals. In "Non-Identity for Non-Humans," they construe Derek Parfit's (1984) non-identity problem for nonhuman animals. The non-identity problem points out

² Among these philosophers are Hanna (2016), Fiet (2015), Bradley (2012), and Feldman (1991).

that I exist due to certain events leading up to my parents meeting; had these not happened, I may not have existed. The same applies to nonhumans too. For example, some birds have had to adjust or change migration patterns due to the effects of climate change.³ These birds would presumably procreate with different mates than they would have otherwise, and consequently, this would lead to a different population of individual birds than would have existed had the climate not been changing.

The collection of individuals in the above example only exist *because of* climate change: this is non-identity for nonhumans. Purves and Hale attempt to show that non-identity for nonhumans threatens what they call “patient-affecting principles.” According to such principles, an act is wrong only if it either harms some moral patient or wrongs some moral patient. But what is harm? Purves and Hale think that patient-affecting principles assume a *counterfactual comparative view of harm*. Proponents of this view think that making an individual or moral patient worse off than they otherwise would have been is harm. For example, suppose my enemy, Jacob, pushes me off of my front porch, breaking my leg. Jacob has decreased my wellbeing. I’m worse off with a broken leg than I am with a good one. Counterfactual comparative indicates that Jacob has harmed me because, had I not been pushed off the porch, I would not have broken my leg. This view compares a possible scenario with the actual scenario and determine in which scenario I am better off.

Let’s clarify counterfactual comparative before proceeding. Consider a slightly modified construal of counterfactual comparative from Erik Carlson:

³ For more information on this topic see Seebacher and Post (2015) “Climate change impacts on animal migration”

Carlson's *Counterfactual Comparative Account of Harm*: An event e or an action a harms [moral patient] S overall if and only if S would have been on balance better off if e had not occurred, or a had not been performed. (Carlson 2018: 2)

For the purposes of this paper, I have tweaked Carlson's rendition of the principle to include *moral patients* rather than *persons*. This modification makes counterfactual comparative applicable to persons and as well as morally concern-worthy non-persons, which are the focus of this paper. Notice further this rendition's focus on *overall* harm. This principle assesses all of the consequences of actions and events and then determines if the individual is overall harmed. Of course, if by saving your life, I accidentally break your leg then I think I have harmed you, but not in an overall sense. Consider another version of counterfactual comparative Carlson entertains:

Maximizing Account: An action a harms a [moral patient] S if and only if there is an alternative action, a' , open to the agent in the situation, such that S would have been better off if the agent had done a' . An action benefits S if and only if there is no such alternative action a' (Carlson 2018: 6)

The maximizing account identifies optimizing possible actions. It is important to note, moreover, that counterfactual comparative is often construed in terms of nearby possible worlds. Justin Kloksiem's version of counterfactual comparative appreciates this relevant difference:

Kloksiem's *Counterfactual Comparative Account of Harm*: a possible event, e (or action, a), would harm [moral patient] S if and only if S is worse off in the nearest relevant possible world in which e occurs (or a

is performed), $W_e(W_a)$, than she is in the possible world nearest $W_e(W_a)$ in which e does not occur (or a is not performed). (Klocksien 2012)

All of these versions of counterfactual comparative are problematic, however. For one, they fail to distinguish harms from failures to benefit. Consider a case:

Suppose that Batman purchases golf clubs with the intention of giving them to Robin, but the Joker persuades him to keep them for himself. Had Batman not kept the clubs he would have given them to Robin (Bradley 2012: 397).

It is clear that Robin is better off with the clubs than without the clubs. According to these accounts of counterfactual comparative, Batman has harmed Robin by not making him better off than he would have been otherwise. Batman appears to be well within his rights to keep the clubs in this case. These versions of counterfactual comparative give us a counterintuitive ruling, a reason to think it misclassifies cases like this. Nathan Hanna (2015) argues that this ruling is not a misclassification at all. Rather, he thinks counterfactual comparative gets cases of failed benefit right. Duncan Purves (2019) disagrees with Hanna. He thinks there is version of counterfactual comparative that distinguishes *allowing* and *making*. With this distinction in mind, he gives us the following counterfactual comparative version of harm:

Harming as Making: An event e is a harm for [moral patient] S if and only if (1) e makes S occupy S 's wellbeing level in the e -world and (2) S 's wellbeing level is higher in the nearest world in which e does not occur. An event e is a failure to benefit S if and only if (3) e does not make S occupy S 's wellbeing level in the e -world, and (4) S 's wellbeing level is higher in the

nearest possible world in which *e* does not occur.
(Purves 2019: 2643-2644)

While this version of counterfactual comparative avoids categorizing failure to benefit as a harm, all of these versions of counterfactual comparative are subject to the non-identity problem. My point here is that the non-identity problem threatens any of these plausible versions of counterfactual comparative. We, therefore, need a response to the non-identity problem.

To return to our discussion, what about the animals who exist because of climate change? Purves and Hale write that if climate change diminishes these animals' welfare, climate change has not harmed them. This claim is puzzling. Why should we think that animals are not worse off? Purves and Hale argue that counterfactual comparative is silent in this case, since we won't find the animals who are worse off in the possible scenarios where the source of "harm" is absent. In possible worlds where the climate is not changing, we see that these animals do not exist there because, as I said above, their existence is contingent upon the changing migration patterns *due to climate change*. Note here that this is only a problem if these animals have lives worth living. If they did not have lives worth living, then it actually would be better that they not exist because their biotic welfare would be zero rather than something negative. So, assuming that these animals have lives worth living, we cannot argue that these animals have been made worse off than they would have been otherwise because otherwise they wouldn't have existed. This, according to Purves and Hale, is a problem for a number of versions of counterfactual comparative.

It is important to notice that patient-affecting principles need not assume counterfactual comparative, so Purves and Hale's argument fails to defeat these principles. For example, as Purves and Hale note, these principles

might assume a “non-comparative” view of harm. Alternatively, if one maintains Nolt’s (2018) view of harm, i.e., an individual I is harmed by an action A only if I is made worse off by some consequence of A than I would have been had that consequence not occurred, we see that many single negative consequences of an action constitutes a harm, such as death, injury, or illness. I will address this account later. My present point is that Purves and Hale’s argument, though it fails to defeat patient-affecting principles, successfully challenges counterfactual comparative. I’ll focus on Purves and Hale’s threat to counterfactual comparative. But before responding to Purves and Hale, we need to get clearer on the concept *species*.

What is a Species?

Species is an ambiguous concept; the definition of *species* is, therefore, controversial among biologists and philosophers of biology. For example, Phillip Kitcher (1984) and Ernest Mayr (1982) distinguish between morphological conceptions of species from genetic conceptions of species. Morphological conceptions of species categorize members into a species based on anatomical features; genetic conceptions, evolutionary lineages and genetic makeup. Ernest argues that while morphology is an indicator of species distinctness, it is not sufficient for distinctness. He writes, “In spite of the variability caused by the genetic uniqueness of every individual, there is a species-specific unity to the... (DNA) of nearly every species” (297). By implication, if we artificially create an orangutan in lab, it may be a member of Bornean orangutan according to morphological accounts but not categorize as a member of B. orangutan on genetic views. This is because this particular orangutan was not part of the same genetic lineage but has the same anatomical features as the rest of the

members of *B. orangutan*. In short, *species* is controversial. My argument is not limited to any one of these conceptions of species; one might think any number of these views are plausible but still find my argument convincing. I will briefly outline a number of views of the definition of species and argue that each one is consistent with my argument.

First, pre-Darwinian, essentialist accounts of species viewed species as natural kinds with unalterable features. Such accounts fail to recognize the various sorts of evolutionary alterations species can go through.

Hull (1978) offers an alternative account of species, in which species are individuals and not classes. Species, according to Hull, are the units of evolution because generations of species are the entities of various hereditary and selection relations. Species are continuous and have spatiotemporal relations. Because classes are not spatiotemporally located, Hull concludes that species are individuals, as opposed to classes. For Hull, the relationship between species and its members is a part whole relation. Ghiselin (1974) warns that "individual" is not synonymous to "biological organism" (573). He thinks there are four features of species as individuals: (1) The species name is a proper name; (2) They do not have defining properties (intensions); (3) There cannot be instances of them; (4) Individual organisms are parts of a species, not members. I think this view is consistent with wellbeing aggregation because the parts of a whole can fare poorly or well. My leg, for example, may not be functional, while my other parts are. At least in principle, we might think welfare aggregation is consistent with this view. I address one version of this view from Holmes Rolston who argues against welfare aggregation.

Kitcher (1984) argues that while some species can be understood as individuals, this is not the case with all species. Biologists use the term *species* in two distinct ways.

First, following Hull, Kitcher suggests that biologists think of species as individuals, which we have seen is consistent with welfare aggregation. Second, biologists sometimes conceive of species as a set of organisms. Kitcher thinks that both senses of *species* are plausible. Both conceptions, I think, are consistent with aggregating welfare. The various members of a set might be doing well. For example, each Koi fish might be doing well in my pond; we then might reasonably say that set of Koi fish is faring well.

Richard Boyd (1999) offers an intuitive account of species that defines a “natural kind” as a “homeostatic property cluster.” According to this view, an individual is a member of a species in virtue of having *many* of the same characteristic as the other members of the species, rather than all of the same characteristics. He thinks that, so long as an individual member of a species has the relevant number of characteristics that the other members have, then the individual can be correctly characterized as a member of that species. This view has problems: first, it is unclear how many characteristics are sufficient for membership in a species; second, if two distinct species share nearly all characteristics in common, then it seems that they should be the same species. There might be other relevant criteria for distinguishing species other than characteristics. In any case, this view is consistent with aggregate welfare because we can look at all of the members of a species and determine how each is faring and then determine if the whole species is faring well.

There are wide range of views available in the literature. It strongly seems to me that each of these views is consistent with present and future welfare aggregation, my view of species wellbeing. I do not think I need to take a stance here on which of these views is the correct one; we should let a thousand flowers bloom.

Species' Harm, Species' Benefit, and Species' Wellbeing

I argue that the wellbeing of a species is the aggregate wellbeing of the present members of the species plus the projection of wellbeing in the near to medium-term future. An endangered species is a species at risk of extinction. They are faring poorly presently such that we can reasonably expect it to fare poorly in the near to medium-term future. But what is faring poorly or faring well? I want to focus on the aggregate *biotic* welfare of the species as opposed to aggregate *experiential* welfare. Biotic welfare is a controversial concept among environmental ethicists.⁴ For the sake of argument, I will not assert any one of these conceptions as the right view of biotic welfare. So long as the view of biotic welfare is aggregable, it will fit with what I am arguing here, and as far as I can tell, any of the mentioned conceptions is consistent with welfare aggregation.

Why merely track biotic welfare of a species, as opposed to just experiential welfare or experiential welfare and biotic welfare? John Nolt (Forthcoming) writes that both dimensions of welfare influence species wellbeing: "Given measures of the biotic and (where relevant) hedonic welfare of individuals, it would be possible to determine an average individual welfare for a species" (7). I disagree.

I have two reasons for thinking we should focus exclusively on biotic welfare. First, tracking just biotic welfare allows us to track species welfare for both sentient and non-sentient organisms. By using just biotic welfare to track species wellbeing, we can track both animals and

⁴ Rolston (1998) argues that biotic welfare is the achievement of normative goals as determined by genetic set, while Nolt (2009) argues that biotic welfare is autopoietic functioning. Nicholas Agar (2001), moreover, thinks biotic welfare is the satisfaction of biotic preferences.

plants. Second, how are we to know what the experiential welfare of an entire species is? I don't think we ever could know that, unless we had a species with only a few members. This concern is augmented when we think about future experiential welfare. How could we ever predict the experiential welfare for future generations? Perhaps we could try, but I do not think we would be very accurate. I am also not sure how prevalent experiential welfare is for species welfare. I suppose that there might be a species that could not reproduce because its members are in a perpetual state of experiential pain but that does not seem plausible. But perhaps we could track experiential value insofar as experiential welfare correlates with biotic welfare: high biotic welfare may track medium to high experiential welfare and low biotic welfare may correlate to low experiential welfare, where we are considering a species with experiential wellbeing. Still, biotic welfare and experiential welfare often come apart. For the sake of simplicity, I merely consider present and projected *biotic* welfare.

Perhaps some may object to this move. Consider that we periodically make claims about the experiential welfare of groups of people, e.g., *Millennials, for the most part, are doing well*. This is a claim, an objector might argue, about the *experiential* welfare of the members of this generation. If this is right, we might think that both biotic welfare and experiential welfare are required for an account of species' wellbeing. The account laid out here, even in light of this consideration, will consider only biotic wellbeing. This is because there are relevant differences between members of a non-human species and members of a generation, e.g., the ability to testify about their welfare, etc.

Species' wellbeing and individual wellbeing are distinct. Something that makes the species worse off does not necessarily make the individual worse off. For example,

consider a species with low population numbers and little genetic diversity. The members of this species could conceivably fare well in the present, but this present wellbeing does not imply that the species fares well. If we project what the species might look like in the coming generations, it strongly seems that welfare would be lower for this species.

A species might benefit from an individual's suffering. Suppose wolves kill and eat a sick and aging caribou. Even though that particular caribou's wellbeing is diminished, its death might benefit the species because its weaker genes do not continue in the gene pool, making for a more advantageous future generation.⁵ This case suggests that the species can benefit in virtue of an individual's harm. But, is this consistent with thinking that species' wellbeing is the aggregated welfare of individuals? I believe so. The individual caribou's suffering lowers the present aggregate welfare of the species, but it (potentially) raises the aggregate welfare of the species in the long term by raising genetic fitness. When we project what the species' population might look like without that individual's genes, we find a slightly stronger, slightly better off species in the future. That projected future good, I think, outweighs the present suffering of the individual. So it's reasonable to think that the species is better off without that individual, even if that individual lowers aggregate welfare in the present.

If a species can benefit in virtue of individual harm, it also seems that the species can be made worse off by individual benefit. Consider overpopulation. Overpopulation might raise the aggregate good for present members of a species, especially if there are enough resources to go around for the present population. Overpopulation,

⁵ This example comes from Nolt (forthcoming) and Rolston (1988).

however, is not a good for the species because it reduces resources for future generations. So overpopulation might be bad for the species but good for the individual and even present aggregate welfare.

These cases show that species' wellbeing and individual wellbeing can come apart although they needn't. We might say that the suffering of an individual is bad for the individual and the species, especially if the species has very few members and if the individual has advantageous genes. We can, therefore, conceive of species' wellbeing as aggregate wellbeing over time. Even if the aggregate welfare of the species is worse off when the caribou is eaten, the species is better off in the long term because of the genetic benefits for future generations.

Holmes Rolston, III (1988, 2012) disagrees; he argues that species' wellbeing is not mere aggregate welfare. In Rolston's view, species are a kind of "super organism" (Nolt, forthcoming). Species do not exist as a class or category of individuals, rather a species is a "corporate individual" as well as a "discrete entit[y] in time as well as space" (Eldredge and Craft 1980: 92). Species as super-individuals can value things, such as "defending a particular form of life, pursuing a pathway through the world, resisting death (extinction), [and] regeneration maintaining a normative identity over time" (Rolston, 151). Rolston thinks that this conception of a species is not compatible with aggregate welfare. He writes:

Duties to a species are not to a class or category, not to an aggregation or average of sentient interest, but to a life line. An ethic about species needs to see how the species *is* a bigger event than the individual interests or sentience. Making this clearer can support a conviction that the species ought to continue. (Rolston 1988: 147)

Even though Rolston's focus is on duties, he thinks aggregating species welfare is problematic. He thinks the wellbeing of a species is something over and above aggregate welfare of the present and future members. My argument depends on the notion that individual wellbeing contributes to the aggregate wellbeing of species; however, if Rolston is right, then aggregate welfare is not the way to assess species wellbeing. Species welfare must be something other than aggregate welfare.

Rolston's view has strengths; it makes sense of the above cases. We might think that the reason species good and individual good come apart is that species good is something over and above individual good. Perhaps a species as a living, historical lifeline has goods that are distinct from the good of its members. Similarly, what is bad for the species might be good for its members; think of the overpopulation example. Rolston's view makes a lot of sense of this asymmetry between species and individual goods. My view also makes sense of this asymmetry because present and future-term aggregate good can look different from individual good. For example, it is bad for the individual elk that it is eaten but benefits the projected aggregate good of future generations.

If we think that a species is an individual as Rolston does, then it is plausible that species have interests. According to this view, a species would presumably have an interest in all the things it can value, such as defending its form of life and pursuing a pathway through the world. Clare Palmer (2011) suggests that one (and perhaps the most) plausible species interest is not becoming extinct. She notes that such an interest is not always obvious though. Suppose that a species in order to continue existing, "all the individual organisms that would compose it, present and future, would have such extremely painful, distressing lives that, as

individuals, they would be better off dead, since their lives are not worth living” (Palmer, 277). As my account of species welfare does not include experiential welfare, we should read Palmer’s concern a low biotic welfare. It is not clear that it would still be in the interest of a species to continue existing. Yet Rolston’s account implies that the species as a “corporate individual” would still have such an interest because the individual has an interest in not becoming extinct over and above its members. This seems implausible in light of Palmer’s concern. Aggregating welfare is the best way to get around Palmer’s worry. If we can reasonably predict that the species will not fare well in the future, then it is possible to say that it no longer has an interest in continuing itself.

If we think about species harm as aggregate welfare of the present population plus future populations, we can make better sense of the non-identity problem’s threat to counterfactual comparative harm relative to climate change. The existence of the *species* is not contingent upon the effects or causes of (anthropogenic) climate change because the emergence of many (perhaps all) present-day species preceded the policies that led to climate change. Thus, we can still talk about climate change’s harm to a species in the counterfactual comparative sense.⁶ The species as a whole might have been better without climate change. I lay out this argument in detail in the next section.

Species’ Wellbeing and the Non-Identity Problem

In this section, I respond to Purves and Hale’s argument that counterfactual comparative fails to make sense of nonhuman suffering. Recalling the non-identity

⁶ Consideration for the *human* species might have some interesting implications for the anthropocentric non-identity problem. Since this issue is beyond the scope of this paper, I merely flag it for future work.

problem, we might think that particular polar bears, for example, might not exist if it were not for climate change.⁷ As the climate changes, polar bears would meet different mates from those they would have otherwise. So, their existence could be contingent upon climate change. Suppose this is true and suppose that climate change diminishes welfare for these polar bears. According to the non-identity problem, we cannot argue that they were harmed in the counterfactual sense because they would not exist without climate change. If we were to think about what would happen “otherwise,” we find most would not exist at all.

Most (or all) species that exist now would have existed even if climate change were not occurring. Even though individual members of a species might owe their existence to climate change or industrialism, the species itself does not. The species of polar bear is much older than climate change and the events leading up to climate change. Suppose now that, due to climate change, the aggregate wellbeing of the polar bear species goes down and the prospective wellbeing of future generations looks low; perhaps the polar bear species becomes endangered such that there are few members and the projection of their future wellbeing is low. Overall biotic welfare is low for these critters. In this case, it is reasonable to think that the species is harmed in the counterfactual sense (on any one of the views I have offered above). Why? We can reasonably talk about what the polar bear species might look like had climate change never happened because there is an *otherwise* to examine. In other

⁷This is the example from Purves and Hale, but I have modified it here slightly. They assert that different polar bears *have actually* come to exist than otherwise had the climate not been changing. This claim is not verifiable and, consequently, implausible. I, therefore, change their argument to a hypothetical in order to bolster it.

“nearby” possible worlds where the climate is not changing, for example, the species exists and strongly seems to be better off. Because the polar bear species’ existence does not depend on climate change, we can think about what it would be like for that species to exist without climate change. In short, we can look at counterfactuals where the polar bear species is better off.

Species’ welfare, I have argued, obtains in virtue of both the welfare of its current individuals and prospective future generations. To use Purves and Hale’s example, suppose some polar bear (whose existence is contingent upon climate change) experiences welfare degradation as an effect of climate change. Let’s call him Knut. Since climate change is responsible for Knut’s existence, following the non-identity problem, we can’t say he was harmed in the counterfactual sense, at least on the views that I have mentioned above. We could, however, look at the effect his suffering has caused to the aggregate welfare of his species. Knut’s welfare degradation probably reduces the overall welfare of the polar bear species, especially if Knut has advantageous genes. Knut’s welfare degradation might also reduce wellbeing of future generations because, perhaps, he is not able reproduce. This would constitute a harm to the species. So even if we cannot say that Knut’s suffering makes him worse off, it is plausible that the species is worse off in virtue of Knut’s suffering on my view of species welfare. And if the species as a whole is better off than they would have been otherwise, then the counterfactual comparative account applies to species wellbeing. Thus, we can really say that Knut’s species was harmed on counterfactual comparative in virtue of Knut’s being a member of the species.

But if it would be better for the species that Knut did not suffer from the harmful effects of climate change, it must

also be better for the species that Knut never existed. In the scenario where climate change does not occur, Knut does not exist so if we say the species is better off without climate change we must also say that the species is better off without Knut. This seems like a problem. But I think this is consistent because the welfare of the species is not contingent on the wellbeing or existence of any particular member. As long as there are more polar bears with greater amounts of welfare in the scenario where Knut does not exist, then it seems right that the species as a whole is better off.

Objections

Objection 1

What about a species whose existence is contingent upon climate change? It's plausible that the changing climate could influence new evolutionary patterns that will occur in the future, such that a new species (who would not otherwise have existed) emerges. Harm done to that species as a result of climate change would fall prey to a collective version of the non-identity problem because this species would otherwise not exist.

I have two responses: first, since this species evolved as a result of climate change, it is also reasonable to think that it will have advantageous features that resist the harmful effects of climate change. This isn't certain, but it is a viable possibility at least. Creatures that are more resilient to climate change have less of a chance of being harmed by it.

Second, in regard to complex mammalian, bird, or reptilian species with long lifespans, this would only happen *way* in the future because the evolutionary process for these critters takes a very long time. This possibility is so far in the future that it probably should not enter our present ethical deliberation. In the event that this does happen, however,

the counterfactual comparative view would be silent about that species' welfare diminishment. That is, if the climate change were to accelerate the evolutionary process and generate new species, then those new species—according to the non-identity problem—would not be harmed by climate change.

Consider a natural rejoinder from my objector: The lifespans of microorganisms are substantially shorter than the lifespans of complex mammalian species. Because of their condensed lifespans, genetic modification and natural selection occur much more quickly than mammalian organisms. Consequently, new species can occur in microorganisms much faster than mammalian species. Here's the problem for my response: I assert that new species will come about much later in the future. But it is likely that new species of microorganisms appear in the near future. Many of these new species may occur as a result of the changing climate, which would result in a kind of non-identity for these microorganisms. Assuming that they are made worse off by climate change, the counterfactual comparative account would be silent about their detriment, i.e., counterfactual comparative could not say that they are harmed.

This objection identifies an authentic weakness in my defense of counterfactual comparative. A microorganism species that exists, in the present or near future, as a result of climate change would not be made worse off by climate harm *per* the non-identity problem. Recall that the non-identity problem undermines a harm assessment in cases where the subject's existence is contingent upon the source of harm. This is because in the counterfactual scenario where the harm does not occur—in this case climate change—that person or species would not exist at all. The question *should we care*, though it lies beyond the scope of this paper, might

offer some way out. For example, the different accounts of counterfactual comparative focus on the betterment or worsening of moral patients. If microorganisms are not moral patients in the relevant sense, perhaps their wellbeing should be of little concern for us. So, maybe it does not really matter that we cannot say these microorganisms are made worse off by climate change. Moreover, to draw from my first response, we might think that these microorganisms have genetic resistance to climate change. Again, this is merely a viable possibility.

Objection 2

An objector could argue that the fact that individual wellbeing and species' wellbeing can come apart is actually a weakness of my account. My argument suggests that we can make sense of counterfactual harm for individuals by looking at harms to the species. But consider the following scenario: suppose only aging and sick polar bears that lack advantageous genetic codes are affected negatively by climate change. If this is true, then (since the death of weak individuals is a good for the species) the species is actually benefitted by climate change substantially. We would expect future generations with greater fitness in this scenario. So, assuming these weaker polar bears' existence is contingent on climate change, the changing climate would not harm any individual according to counterfactual comparative. We cannot make sense of individual harm by looking at the species in this case, given the benefits to the species.

Response: Though this scenario points to a weakness in my defense of counterfactual comparative, I am not sure climate change actually works this way. It appears that climate change could harm any number of fit individuals in addition to weaker, less fit individuals. If climate change affects both weak individuals as well as fit individuals, then

it seems like the species would still be worse off. Thus, we can track individual harm with species harm and use counterfactual comparative to track the harm.

Besides, even if it is sometimes beneficial for a species to have weaker members die, surely it is not always beneficial. Suppose that climate change kills *all* weaker members of a given species, such that there are few fit individuals left. This would clearly make the species worse off, even though only weaker members of the species are affected. Thus, even though most of the time the death of weaker individuals is beneficial, it does not follow that this is true in all cases.

Objection 3

Consider yet another objection. There's a version of counterfactual comparative that the non-identity problem does not threaten, which Atkins (2018) has argued in favor of. John Nolt (2018) offers another version of counterfactual comparative. He thinks that so long as at least one of the consequences of an action makes an entity worse off, this counts as harm: "People are harmed (in a comparative sense) by an action or policy only if at least one of its consequences makes them worse off than they would have been had that consequence not occurred" (5). Consider what I am calling principle H:

Principle H: an individual I is harmed by an action A only if I is made worse off by some consequence of A than I would have been had that consequence not occurred.

Note that principle H offers a necessary condition of harm. The strength of this account is that it leaves open the possibility of simultaneous harm and benefit. Rather than aggregating *pro tanto* harms and *pro tanto* benefits to

determine if a moral patient is harmed or benefited *overall*, Nolt thinks an action is harmful or beneficial so long as it has at least one beneficial or harmful consequence. But first, what is a *pro tanto* harm as opposed to an *overall* harm? A *pro tanto* harm is a specific consequence of an action that must be considered alongside other *pro tanto* harms and *pro tanto* benefits. *Pro tanto* harms can be outweighed by a sufficient number of *pro tanto* benefits and vice versa. In a case where I break your leg to save your life, I would, according to this view, both harm and benefit you rather than benefit you overall.

On this view, we can make sense of the non-identity problem. For example, climate change in some cases harms and benefits those who exist in virtue of climate change. We would not need to look to an alternative state of affairs to see if a moral patient is harmed; rather, we examine each individual consequence of climate change to see if harm has been done. Presumably, if your life is worth living, you're better off existing (if you exist because of climate change), even if climate change *pro tanto* harms you. The non-identity problem is only a problem when we aggregate *pro tanto* harms and benefits to determine if a moral patient is better or worse off overall; this version of counterfactual comparative avoids aggregation and, consequently, the non-identity problem. We can, therefore, make sense of the non-identity problem *and* keep at least one version of counterfactual comparative. The argument of this paper, therefore, misses this important reconciliation of the non-identity problem and counterfactual comparative; my argument isn't necessary.

In response, Nolt's account of harm doesn't take seriously the possibility of being harmed in an *overall* or all-things-considered sense. The counterfactual comparative account needs the distinction between *overall* harm and *pro*

tanto harm. A *pro tanto* harm is a single consequence of an action that might be outweighed by a set of *pro tanto* benefits that are consequences of a given action. So a person can experience a *pro tanto* harm but still not be worse off *overall*. In addition, we cannot make sense of certain statements without this distinction. Consider, for example, the following news headline from Bradley (2012): New studies show surgery is harmful! There are plenty of *pro tanto* harmful consequences of surgery: pain, bodily dismemberment, cutting, etc. These all make the person worse off to some degree. It would be odd, however, if all this statement referred to were these *pro tanto* harms. After all, we expect these sorts of harms after a surgery, since they are obvious consequences of many surgeries. The headline would be trivial if it refers *only* to *pro tanto* harms. This statement is interesting and meaningful only if it asserts that surgery makes you worse off *overall*. Suppose that new data emerges that suggests certain surgeries shorten life expectancy. Even though your ailment is cured, you cannot expect to live very long. This new finding would only make sense if we consider harming *overall*. Because Nolt's account fails to distinguish between overall and *pro tanto* harm, it cannot make sense of the above statement.

Let's construe principle H with the distinction between overall and *pro tanto* harm. I'll argue that one version of the account is false and the other is obvious:

H*: an individual I is *overall* harmed by an action A only if I is made worse off by some consequence of A than I would have been had that consequence not occurred.

This reading of H is false. Consider again the surgery case. Even if I experience some minor pain briefly after the surgery, I would not be worse off overall because, say, I

would have died if I had not had the surgery. Now let's think about a second reading of the principle:

H^{**}: an individual I is *pro tanto* harmed by an action A only if I is made worse off by some consequence of A than I would have been had that consequence not occurred.

H^{**} is clearly true. This construal of the principle does not, however, have much bite; it's obvious. Again, consider the surgery case. It seems odd to think that surgery harms me, even if there is some minor welfare reducing outcome.

Could Nolt reject this distinction? Perhaps every *pro tanto* harm is a harm. This response is problematic because by rejecting this distinction, we significantly broaden the harmful too widely. Broadening harm this widely is counterintuitive. *Anything* with at least one welfare-reducing consequence would count as harmful. Thus, a surgery, while beneficial in many ways, is harmful according to this account so long as there is one welfare-reducing consequence. We would have to say that surgery is both harmful and beneficial. Again, we could not make much sense of the headline I mentioned above because it makes the most sense under an all-things-considered account of harm.

On Nolt's account, it seems that every surgery is harmful simply in virtue of having *at least one* welfare-reducing consequence. But it seems wrong to conclude that surgery *just is* harmful. We need to aggregate the number of harmful consequences and beneficial consequences and then determine if there is overall more harm or more benefit. But when we do that, I think that the appeal of this account diminishes.

Conclusion

In this paper, I have argued that the non-identity problem does not threaten counterfactual consideration of species welfare. Because entire species existed before climate change, the existence of the species is not contingent on it. It's, therefore, plausible to use the counterfactual comparative account of harm to assess the harm done to an entire species. Although individual welfare degradations are subject to the non-identity problem, I have argued that individual welfare degradations are harm to the species.

Bibliography

- Atkins, J. 2018. "Have You Benefited from Carbon Emissions? You May Be a 'Morally Objectionable Free Rider.'" *Environmental Ethics*. 40 (3): 283-296.
- Bradley, B. 2012. "Doing Away with Harm." *Journal of Philosophy and Phenomenological Research*. 85 (2) 390-412.
- Boyd, R. 1999a, "Homeostasis, species, and higher taxa", in R. Wilson (ed.), *Species: New Interdisciplinary Essays*, 141-185. Cambridge: MIT Press.
- Carlson, E. 2018, "More Problems for the Counterfactual Comparative Account of Harm and Benefit." *Ethical Theory and Moral Practice*. 22 (4): 795-807.
- Eldredge, N. and J. Cracraft, 1980. *Phylogenetic Patterns and the Evolutionary Process*, Columbia University Press.
- Feit, N. 2015. "Plural Harm." *Philosophy and Phenomenological Research* 90 (2): 361-388. doi:10.1111/phpr.12033.

- Feldman, F. 1991. "Some Puzzles About the Evil of Death." *Philosophical Review* 100 (2): 205-227. doi:10.2307/2185300.
- Ghiselin, M. 1974. "A Radical Solution to the Species Problem." *Systematic Zoology* 23: 536-44.
- Hanna, N. 2016. "Harm: Omission, Preemption, Freedom." *Philosophy and Phenomenological Research* 93 (2): 251-273. doi:10.1111/phpr.12244.
- Hull, D. 1978, "A Matter of Individuality", *Philosophy of Science*, 45: 335-360.
- Kitcher, P. 1984, "Species", *Philosophy of Science*, 51: 308-333.
- Klocksien, J. 2012. "A Defense of the Counterfactual Comparative Account of Harm." *American Philosophical Quarterly*. 49:285-300.
- Mayr, E. 1982. *The Growth of Biological Thought: Diversity, Evolution, and Inheritance*. Belknap Press of Harvard University Press.
- McMahan, J. 2013. "Causing People to Exist and Saving People's Lives." *Journal of Ethics*, 17 5-35.
- Nolt, J. (Forthcoming). "Climate Change and the Welfare of Life." In *Climate Change Ethics and the Non-human World*, Brian G. Henning and Zack Walsh, (eds.). Routledge.
- . 2018. "Long-term Climate Justice." In *Climate Justice: Integrating Economics and Philosophy*, Ravi Kanbur and Henry Shue, eds., Oxford: 230-246.
- Palmer, C. 2011. "Does nature matter? The place of the nonhuman in the ethics of climate change." In *The Ethics of Global Climate Change*. Denis Arnold (ed.) Cambridge: Cambridge University Press, 272-291.

Parfit, D. 1984. *Reasons and Persons*, Oxford University Press.

Purves, D. 2019. "Harm as making worse off." *Philosophical Studies*, 176: 2629-2656.

Purves, D. and B. Hale, 2016. "Non-Identity for Non-Humans," *Ethical Theory and Moral Practice* 19:1165-1185, doi 10.1007/s10677-016-9752-3

Rolston, III, H. 1988. *Environmental Ethics: Duties to and Values in the Natural World*, Temple University Press.

———. *A New Environmental Ethics: The Next Millennium for Life on Earth*, Routledge.

Seebacher, F., Post, E. 2015. "Climate change impacts on animal migration." *Climate Change Responses* 2, 5 doi:10.1186/s40665-015-0013-9