Climate change, intergenerational justice, and the non-identity effect

Bontly, Thomas D.

Veröffentlichungsversion / Published Version
Zeitschriftenartikel / journal article

Climate change, intergenerational justice, and the non-identity effect

by Thomas D. Bontly

Abstract: Do we owe it to future generations, as a requirement of justice, to take action to mitigate anthropogenic climate change? This paper examines the implications of Derek Parfit's notorious non-identity problem for that question. An argument from Jörg Tremmel that the non-identity effect of climate policy is "insignificant" is examined and found wanting, and a contrastive, difference-making approach for comparing different choices' non-identity effects is developed. Using the approach, it is argued that the non-identity effect of a given policy response to climate change depends on the contrasting policy. Compared to a baseline scenario without further mitigation, the non-identity effect of choosing to limit climate change to 1.5°C would be highly significant.

Keywords: Climate change, Non-identity problem, Intergenerational justice

Introduction

Many of the questions we confront today have profound implications for the lives and living conditions of future people. Should we reduce carbon emissions? Conserve resources? Pay down the national debt? Curb population growth? Whatever the case may be, it's hard to believe that we could respectably address such issues without paying some attention to interests of future people.

And yet, when we think about justice for future generations, we run up immediately against Parfit's (1984) notorious non-identity problem. In this paper, I take a fresh look at the problem's implications (or lack thereof) for policy responses to anthropogenic climate change. Recently, Tremmel (2018) has argued that the non-identity problem only arises in unrealistic, "fact-insensitive" thought experiments; in the real world, where climate change is an urgent problem, the non-identity effect of our policy choices is "insignificant" (2018: 44) and can be safely ignored. Here I develop a difference-making approach for thinking about the significance of such effects and subsequently argue that the non-identity effects of climate policy are indeed significant. If one wishes to avoid the non-identity problem, one will have to tackle it head on.

The non-identity effect

In what follows, it will help to distinguish the non-identity *problem* itself from what I call, following Broome (2012: 62), the non-identity *effect*. A choice has a non-identity *effect* if it makes a difference to who subsequently comes to exist. Choosing to have one child rather than none has a non-identity effect, as does choosing to have a child with one mate rather than another or even with the same mate at a different time.

The non-identity *problem*, on the other hand, arises when a choice strikes us as morally objectionable, on account of its effects on some person, even though (due to its non-identity effect) that very person would never exist if the choice went another way. For instance,

> *Zika*. A couple living in a region where Zika virus is circulating wishes to have a child. Their doctor advises them to wait a month before conceiving, by which time the risk of infection will have passed. But they are in a hurry and conceive forthwith. The woman is then bitten by a Zika-carrying mosquito; the infection is transmitted to the foetus; the child is born with microcephaly and has reduced quality of life as a result.

Was their choice to conceive forthwith morally objectionable? One tends to think so. That is not to say their choice was wrong, all things considered; perhaps they had good reason not to delay. Still, they had a moral reason to delay, and presumably that reason had to do with the welfare of their child. Their child would have had a better life, one supposes, had they elected to wait.

A choice has a non-identity *effect* if it makes a difference to who subsequently comes to exist. The non-identity *problem*, on the other hand, arises when a choice strikes us as morally objectionable, on account of its effects on some person, even though (due to its non-identity effect) that very person would never exist if the choice went another way.

The problem, of course, is that it would have been a different child having that better life. Had they waited, their actual child would never exist at all. *That* child is no worse off than she would have been had they waited, because her existence, however impaired, is not worse than no existence at all.[1] Nor has that child been harmed, if harming someone requires making her worse off. Consequently, it is difficult to explain what our objection to the couple's choice might be.

But the non-identity effect is not restricted to procreative choices, as Parfit (1976) first pointed out. Indirectly, socio-economic policies have non-identity effects as well, by affecting our lives in countless ways – where and how we live, work, study, play – thereby affecting whether, when, and with whom we have children. Which brings us back to the subject of climate change. Climate change is, as Gardiner (2006) puts it, a severely time-lagged phenomenon, the effects of which are heavily backloaded. The changes we see now – and the consequent casualties[2] – result from CO_2 emissions accumulated over the last two centuries. The impact of current and future emissions, on the other hand, will be felt some time (and, due to the long residence of CO_2 in the atmosphere, for a *long* time) in the future. Most of those impacts will fall on people not yet born.

Based on current projections, furthermore, the impacts are expected to be disastrous. Absent significant mitigation,[3] the Intergovernmental Panel on Climate Change (IPCC) projects global mean surface temperature will rise 3.7°C to 4.8°C over preindus-

trial levels by 2100.[4] Warming of that magnitude would bring "high to very high risk of severe, wide-spread and irreversible impacts globally" (IPCC 2014a: 19), including the loss of species and ecosystems, extreme weather events, significant and irreversible sea-level rise, and, for humans, increased food and water insecurity, disease, dislocation, conflict, and poverty.[5] The good news is that it is still possible to reduce climate change risks through concerted efforts at mitigation (IPCC 2014c:14). The bad news is that doing so will be extremely expensive; no doubt many other projects and opportunities would have to be sacrificed.[6]

Hence the question: do we owe it to future people, as a requirement of justice, to act now to prevent these bad effects from occurring? Many think so.[7] But now we must confront the non-identity problem. Since the choice to mitigate climate change (or not) will affect who later exists, we cannot claim that future people will be individually better off if we reduce our emissions. *Other* future people would enjoy a more stable climate and better living conditions if we reduce emissions, but not *the same* future people.[8] Nor will future individuals be able to claim, if we fail to mitigate climate change, that they had a right to inherit a better world, for we could not have left *those* future people a better world (Broome 2012: 62). How then can it be maintained that we have an obligation to future people to act?

It seems, therefore, that the non-identity problem has the potential to undermine justice-based arguments for mitigation. For similar reasons, it may undermine arguments of historic justice that inhabitants of industrialised countries, having benefited from past industrialisation, ought now to bear the burdens of mitigating and adapting to climate change. Thanks to the non-identity effect, however, it is not true that inhabitants of industrialised countries are better off than they would have been had industrialisation not occurred; without industrialisation, *those* individuals would not exist (Caney 2005: 757-758).

Can we dodge the problem?

Thus far, we have seen that the non-identity problem has the potential to undermine two familiar arguments about climate justice. Of course, if it could be shown that the non-identity problem rested on some error, those arguments might hold up rather well. I shall return to that possibility below. First, though, I would like to consider three attempts to sidestep the non-identity problem by arguing that its implications for policy in this area are quite limited.

One way to evade the non-identity problem would be to appeal to moral obligations which are not duties of justice. According to Broome (2012: 52-53), duties of justice are owed to particular people, whereas duties of goodness are owed to no one in particular. Our governments, furthermore, have a general obligation of goodness "to promote the flourishing of their people" (65), giving them a reason, quite apart from justice, to mitigate climate change. However, shifting our focus from justice to beneficence does not render the non-identity problem moot. First, when justice and beneficence conflict, justice usually takes priority. If we owe it to current people to help them adapt to climate change, the idea of mitigating climate change in order to promote future flourishing would have to take a backseat; whereas if we owe it to future people to mitigate climate change as a duty of justice, their claim may win out. Furthermore, our governments are apt to interpret the duty to promote the flourishing of "their people" rather narrowly, to include only those who might vote for them in the next election. Promoting the good of future people, while nice, seems supererogatory – unless, of course, one owes them something as a duty of justice.

A second attempt to evade the non-identity problem might point out that the negative effects of climate change are not deferred as far into the future as scientists used to think. Climate change is already killing people and will almost certainly make life increasingly miserable for a great many who already exist.[9] Today's children, after all, can reasonably expect to see the year 2100, which is the endpoint for most IPCC projections. Since those children already exist, our duty to mitigate climate change for their sake is not undermined by the non-identity effect, which only affects duties to future people.

A second attempt to evade the non-identity problem might point out that the negative effects of climate change are not deferred as far into the future as scientists used to think. Let us grant that we owe it to our existing children to mitigate climate change. Still, the non-identity problem is not a moot point as it bears on the need to weigh the costs and benefits of mitigating climate change against those of adapting to it.

Let us grant that we owe it to our existing children to mitigate climate change. Still, the non-identity problem is not a moot point as it bears on the need to weigh the costs and benefits of mitigating climate change against those of adapting to it. As Moellendorf (2015: 174) observes, adaptation policies can benefit those alive now as well as future people, whereas mitigation mostly benefits future generations – unless, that is, the non-identity effect means that we cannot benefit future people at all. And if mitigation would benefit practically no one, surely we ought to direct our resources toward adaptation instead. So, even if we have ample reason to worry about climate change for the sake of current people, the non-identity effect still has implications for the appropriate policy response.

Is the effect insignificant?

A third and I think more interesting reason why the non-identity effect might be irrelevant to climate policy is suggested by Tremmel (2018). As he sees it, the non-identity effect only matters in unrealistic, "fact-insensitive" thought experiments where the causal factors at play are artificially expanded. In the real world, where climate change is an urgent problem, the non-identity effect of policy choice is "miniscule" and "insignificant" (44) and can be safely ignored.

Why insignificant? As Tremmel points out, government policies aimed at controlling emissions aren't the only factors affecting whether, when, and with whom people make babies – far from it. A myriad factors play a role: college admission policies, dating apps, and the closing times of bars play, in his view, just as big a role (2018: 46), to say nothing of such things as trade policy, tax incentives, housing prices, war, financial deregulation, and on and on. But the more factors there are, the less any one seems to matter: "If the number of factors that influence who will be meeting, mating and making children with whom converges toward infinity, the influence of each particular factor converges towards zero" (46).

Let us grant that a great many factors affect who later comes into existence. It does not follow, however, that the effect of any one such factor is small or insignificant. We can see this from classical physics. According to Newton's laws, the acceleration a particle undergoes depends on its mass and on the force applied to it, and every particle with mass exerts some gravitational force on every other. Thus, the number of factors influencing the motion of any one is potentially infinite. Does it follow that the influence of each must approach zero? No. Since gravitational force is proportional to mass and inversely proportional to the square of distance, the influence of nearby, massive particles is bound to be significant by comparison to the rest, no matter how many there are.

While that suggests that Tremmel's argument is unsound, classical physics is a poor model for the non-identity effect. Physical forces obey Mill's (1858) principle of the composition of causes: the result of two forces is just the sum of each were it acting alone. Consequently, we can ask "How much of A's acceleration is due to B?", because force is a quantity, and quantities can be aggregated. It's like asking how much of one's martini is gin and how much vermouth; forces, like spirits, can be aggregated.

But the non-identity effect does not work that way. The question "How much of the population of the future is due to climate policy?" makes little sense, because policies don't contribute discrete sets of individuals; nor do dating apps, bar times, or banking regulations. We cannot coherently ask who would exist in the future if climate policy (or dating apps, etc.) were the only "force" at work, for such a situation is inconceivable.

So, Tremmel's attempt to sidestep the non-identity problem seems unworkable, but the analogy with physics suggests a new question: how should we think about the contributions of different identity-affecting factors?

A difference-making account of causal significance

For our purposes, a better model than physics is biology. Sometimes we want to ask whether a given trait, e.g. height, depends more on the organism's genes than on its environment. But we do not suppose that genes and environment make separate contributions which can be aggregated like forces (Sober 1988). It makes no sense to ask how tall someone would have been if genes had acted alone, or if environment had acted alone. Genes cannot act without environment, nor environment without genes; both are required for phenotypic effects (Ariew 1996).

Still, there are ways to compare the effects of genetic and environmental factors. One way is to ask, of a trait in a given population, whether genes or environment make more of a difference. That is the intuitive idea behind the statistical technique known as *analysis of variance* (ANOVA). To illustrate briefly, suppose we have a number of plant seedlings representing three genotypes (G1, G2, G3) of the same species.[10] We plant them in separate plots, varying only the amount of water they receive (low, medium, high). We then measure their heights at maturity, average them, and plot as a function of the other two variables, as in Table 1.

	W=High	***W=Medium***	***W=Low***	***Marginal average***
G1	85	55	25	55
G2	80	50	20	50
G3	75	45	15	45
Marginal average	***80***	***50***	***20***	
				Grande mean = 50

Table 1: Fictional data for plant heights and water-breed combinations

The variance, intuitively, is the extent to which the values in the nine cells differ from the grand mean. The question ANOVA answers is: which variable accounts for more of the variance? (In this case, the answer is W.)

Four points about the statistical analysis of variance are relevant to our discussion. First, which factor accounts for more of the variance is relative to a population. In this population, W's value makes a bigger difference to height than G's. In a population containing other variants (say, G3, G4, and G5), G(enotype) might account for more variance than W(ater).

Second, statistical difference-making is *contrastive*. The difference a given value of G (say, G1) makes to H depends on what other value(s) of G we contrast it with, or what we take as the "baseline". G's being G1 rather than G2 makes a difference of 5 units, whereas G's being G1 rather than G3 makes a value of 10 units. So, our choice of contrast matters.

Third, nothing here turns on how many different factors are at play. Presumably, a myriad of factors, both environmental and genetic, influence height; even so, it would not follow that water's effect on height is insignificant (though we might find it to be so in some populations).

Fourth, ANOVA is a population-level analysis, inapplicable to a singleton case. In a population of one, there is no variance to analyse. And that is a problem for us, because in non-identity cases, we really are dealing with a population of one – not of one person, necessarily, but a population of one population. There is only one human population, and we are wondering about the effect climate policy has on its composition. If we had a bunch of populations, all genetically identical, we could "plant" them (so to speak) and subject them to various manipulations. But we do not.

Still, we can pretend. Imagine, as a variant on Putnam's (1975) Twin Earth case, that we had the god-like power to make hundreds of Twin Earths, each a molecule-for-molecule duplicate of the Earth as it is right now, right down to the DNA molecules of all 7 billion of us. Then we divide the many Twin Earths into distinct lots, subjecting each lot to a distinct set of manipulations. Some Twin Earths we subject to a mix of policies aimed at keeping warming by 2100 below 1.5°C; others we subject to a mix of policies aimed at keeping warming below 2°C; others receive no additional mitigation. We can manipulate other factors as well, e.g., bar times. Then, at various points in the future (say, 30-year intervals), we check to see if the "same people" (genetic twins) are born.[11]

Obviously, we cannot do any of these things, except in thought, which brings me to a fifth point.

In the singleton case, the relevant notion of difference-making is not statistical; it is counterfactual (Sober 1988). To what extent is Giorgione's height due to his environment? The only way I can see to answer that question is with a counterfactual: *to the extent that his height would have been different had he the same genes but been raised in a different environment*. As with statistical differ-

ence-making, furthermore, counterfactual difference-making is contrastive: that is, we must specify in what other environment(s) Giorgione would have been raised if not his actual. If he was raised in a high-nutrition environment, the interesting contrast might be a low-nutrition environment, in which case (perhaps) he would have been far shorter; then his environment made the difference. Or perhaps he would have been just as tall, in which case his environment made little difference. Whichever the case, contrastivity seems unavoidable.

In the singleton case, the relevant notion of difference-making is not statistical; it is counterfactual. Whichever the case, contrastivity seems unavoidable.

Why the non-identity effect of climate policy is significant

Now, I suggest that we frame our question about the non-identity effect of climate policy in the same terms. Consider the most ambitious goal set forth in the 2015 Paris Agreement, that of limiting the increase in global average temperature to no more than 1.5°C, and suppose that we have some idea of the policy regime (a mix of carbon taxes, infrastructure investments, and so on) needed to achieve it. How significant would the non-identity effect be? What difference would it make to "the phonebook of the future" (in Tremmel's apt metaphor)?

It depends, I submit, on what policy(s) we select as the alternative. One alternative is the "business as usual" or "baseline" scenario involving no mitigation measures beyond those in place; another is the slightly less ambitious goal highlighted in the Paris Agreement of holding warming by 2100 below 2°C. While 0.5°C does not sound like much difference, IPCC (2018) finds major differences in what it would take to achieve them. Based on current models, the 2°C goal would require emissions to decline 25% by 2030 and reach net zero by 2070, whereas 1.5°C would require emissions to decline 45% by 2030 and reach net zero by 2050. As IPCC (2018: 15) puts it, in comparison to 2°C, limiting warming to 1.5°C "would require rapid and far-reaching transitions in energy, land, urban and infrastructure (including transport and buildings), and industrial systems" on an unprecedented scale and "imply deep emissions reductions in all sectors, a wide portfolio of mitigation options and a significant upscaling of investments in those options".[12]

Now, the first point I want to make is that the difference between these "rapid and far-reaching transitions" required for 1.5°C and the somewhat less rapid but still far-reaching transitions required for the 2°C goal would make one sort of difference to future population. By contrast, the difference between the 1.5°C scenario *and the baseline* would be quite another. Of course, we cannot say exactly how much the future population would differ in either case, much less which individuals would exist. But there is reason to think that, compared to a policy of business-as-usual, the adoption of policies consistent with 1.5°C of warming would significantly alter the details of most people's lives. If so, their future populations would in all likelihood diverge rapidly, with entirely different people being born in fairly short order.

There are several reasons to expect rapid divergence. First, economists agree that any serious attempt to reduce emissions would require putting a price on carbon emissions; the steeper the target reduction, the higher the price would have to be initially and the more rapidly it would have to increase. Under a less ambitious target, we might limit the impact of carbon pricing on people's lives by rebating some or all of the revenues back to taxpayers, or by using it to address poverty. But if we need to reduce emissions by 45% by 2030 and reach net zero by 2050, it's plausible that most if not all of the revenue would have to flow towards infrastructure, research and development. Compared to baseline, therefore, those in a 1.5°C world would for all intents and purposes live in vastly different economies; they would attend different schools, enter different careers, live in different places, travel by different means, and meet different people, all of which would subtly affect the timings of conceptions.

A further difference between the baseline and 1.5°C target scenarios involves climate change itself, which has add-on non-identity effects. Under the 1.5°C target, the climate will warm less rapidly and less overall compared to baseline, reducing both risk and the cost of adaptation in decades to come. There would be many fewer climate refugees than under baseline. These differences would further affect people's lives and indirectly their procreative choices.

A third difference between the baseline and 1.5°C target scenarios involves social unrest. We have already seen, in the Yellow Jacket movement in France, how one country's quite modest attempts to curtail emissions by raising fuel taxes can lead to unrest, and unrest has non-identity effects of its own. Unrest brings protesters together, but it can also drive neighbours apart. It brings some into the street while leading others to stay home, all of which affects who meets and ultimately mates with whom and when. Under the baseline scenario, too, we should expect social unrest, though in different communities and with different results. Under baseline, the protesters would be primarily young people and progressives (in affluent countries in the Global North); under 1.5°C, the protesters are more likely to be older middle-class men unhappy with the added cost of living.

A fourth difference between baseline and 1.5°C scenarios is a difference in our values. Barring technological miracles, rapid decarbonisation may not be achievable in capitalist democracies, where economic growth is normative, multinational corporations hold immense power, economic inequality runs high, and the accumulation of wealth and material goods are employed as the measure of a life. To rapidly transform society, therefore, we likely must simultaneously transform our media, our schools, our politics, and our values. And since values guide choices, our children will likely make different procreative choices, compared to baseline, in part because of their different values.

Barring technological miracles, rapid decarbonisation may not be achievable in capitalist democracies, where economic growth is normative, multinational corporations hold immense power, economic inequality runs high, and the accumulation of wealth and material goods are employed as the measure of a life.

Compared to baseline, therefore, one should expect the adoption of policies aimed at 1.5°C of warming to have a large non-identity effect; due to the far-reaching changes needed in society, quite possibly no one's life would be quite the same. In terms of the analogy to classical physics from above, adopting such policies would be like adding a distant but extremely massive object to our solar system: the effect on other bodies would be relatively indi-

rect, but it would subtly affect the trajectories of them all. Compared on the other hand to policies aimed at 2°C, the non-identity effect of 1.5°C would be smaller but still, due to the shorter time to reach net zero, quite pronounced.

The outlook

In this paper, I have argued that we cannot avoid the non-identity problem in climate ethics by supposing that the non-identity effects of our various policy options would be insignificant or otherwise beside the point. Hence, if we are persuaded that unmitigated CO_2 emissions would do future people an injustice, we must attack the problem head on: we shall have to show that an act can harm or otherwise wrong someone who would never exist but for that very act.

While I lack the space to defend such an approach, I shall briefly argue that one of the assumptions underlying the problem is quite vulnerable and suggest an improvement. This is the assumption that *an act cannot harm someone if it does not make her worse off than she would otherwise have been*. Although there is a connection between harming and counterfactual difference-making, it is not as simple as the italicised formula would indicate. We can see this by considering a different thought experiment, one involving causal preemption:

Thirsty Traveller. A traveller, T, sets out on a trip across the desert. T has two enemies, A and B. A puts a deadly poison in T's reserve can of drinking water. Then B, unaware of the poison, drills a hole in the bottom of the can. By the time T needs the reserve, the can is dry; T dies in the desert.[13]

In this case, *someone* harmed T, and it certainly wasn't A. After all, A's poison never touches T's lips; his attempt on T's life is cut short. We can agree that A wronged T by *trying* to harm him, but it also seems clear that B wrongs T in a more direct sense – by actually harming him. And if B harms T, then harming someone cannot require making that person worse off. For in this case, T winds up no worse off than he would have been if B hadn't drilled his little hole (Bontly 2016: 1237).

Preemption cases suggest that we need to rethink the counterfactual account of harm. As a first pass, consider a simple (though ultimately inadequate) causal account of harm, where an action harms someone if it actually causes something to occur that is worse for that person, i.e. if it produces an effect that person would be better off without. In *Thirsty Traveller*, due to the preempted backup, B's drilling the hole does not make T worse off than T would otherwise have been. However, B's drilling does *cause* something – viz. T's death – which leaves T worse off than he would otherwise have been. Insofar as preemptive harms are concerned, thus, the simple causal account appears to improve upon the familiar counterfactual account of harm.

Similarly, in the case of climate change, the simple causal account appears to vindicate the commonsense view that our choices can harm future people, even those who would never exist but for those choices. To keep things manageable, let us focus on a dichotomous choice: either to pursue mitigation policies sufficient to limit warming to 1.5°C (henceforth, "Mitigation"), or to continue our reliance upon fossil fuels without any attempt to mitigate ("Baseline"). Suppose now that we choose Baseline, extreme warming ensues as predicted, and millions of people in the 22nd century suffer or die prematurely from climate-related causes. Let us assume, furthermore, that none of these millions of people would ever exist if we chose Mitigation, due to the non-identity effect, though other people would. A simple counterfactual theory of harm tells us, counterintuitively, that Baseline does not harm those future people, for they themselves are no worse off than they would be under Mitigation. The causal account, on the other hand, says just the opposite: our choice harms the future people, because it causes extreme warming, which is worse for them than the lesser warming under Mitigation would have been.

So, a causal account of harm can explain why there is at least a *pro tanto* objection to the Baseline choice, that is, why it is objectionable in some respect or to some extent. By that same token, however, it may seem on a causal account that the *pro tanto* objection is overridden by a greater *benefit* we give to those same people: namely, the benefit of existence. For we may assume that existence is on balance good for these future people – that the good they experience in their lifetimes outweighs the bad, despite the ill effects of unmitigated climate change. And, of course, these future people would receive none of these benefits if we had not chosen Baseline, for then they would not exist at all. Thus, it can be argued that, on a causal account, our choice benefits the future people more than it harms them, leaving it unclear why that choice is objectionable, all things considered.

However, a subtle amendment to the causal account solves the problem. On the view I favour, harming and benefiting are not just causal notions; they are *contrastive* notions. That is, the claim that some action, x, harms a particular person must be understood as the claim that the performance of x *rather than some specific alternative(s) x** harms that person, where x* is the act (or set of acts) the agent would or might have done instead. Then the contrastive account of harming (and benefiting) runs as follows:

The performance of x rather than x* harms (benefits) someone, S, if and only if (i) there are events e and e* such that x rather than x* causes e rather than e*, and (ii) e is worse (better) for S than e* would have been (Bontly 2016: 1246-1247).

Now consider the choice between Baseline and Mitigation. Just let e be extreme warming of (say) 4°C or more, and let e* be warming of only 1.5°C. By hypothesis, 4°C is worse for those people than 1.5°C would be, and Baseline rather than Mitigation causes warming of 4°C rather than 1.5°C. So, our choice does indeed harm the future people.

Furthermore, there is no reason, on the contrastive account, to think that our choice benefits the future people in causing them to exist. For it to benefit them, there would have to be some e and e* such that (i) Baseline rather than Mitigation causes e rather than e*, and (ii) e is better for the future people than e* would have been. One can easily find pairs of events that satisfy one constraint or the other, but no events in our scenario seem to satisfy both conditions simultaneously. For instance, our choosing Baseline over Mitigation causes one group of future people to come to exist rather than another. But is it better for these future people that they exist rather than some other future people? No. A person isn't better off in worlds where she exists than in worlds where she never exists at all. Nor is she worse off or equally well off. Such comparisons presuppose that the person has a welfare in both worlds, and that presupposition is not satisfied in this case.

On the other hand, assuming that our future people have lives that are on balance good ones, there must surely be events in their lives which satisfy (ii). The question is whether our choice causes such events to occur rather than alternatives that would be worse

for them. Suppose that one of these future people is a faculty member, and that her receiving tenure is better for her than her being denied tenure would have been. But choosing Baseline over Mitigation does not cause our future professor to receive tenure rather than to be denied tenure, nor does it clearly cause her to have any other good rather than to lack it (Bontly 2016).
These claims require further defence than given here, but I am hopeful that a contrastive account of harm will deliver us from the non-identity problem.

Notes

1 Nor, presumably, is their actual child better off than she would have been had they waited. A person is better off in scenario A than in scenario B only if her wellbeing in A would be higher than in B. Since a person has no wellbeing in worlds where she has no being, one cannot (I conclude) be better or worse off than if one never existed. For a contrasting view, see Roberts (2003).
2 Estimates of deaths from climate change vary. One oft-cited number comes from the World Health Organization (2009: 24), which estimated that climate change was already responsible for 140,000 excess deaths annually by the year 2004.
3 "Mitigation" in this context refers to reducing CO_2 emissions or enhancing CO_2 sinks in order to stabilise the climate; "adaptation" means reducing human vulnerability to risks from a changing climate.
4 The range of 3.7°C to 4.8°C assumes a median estimate of climate sensitivity to increased levels of greenhouse gas. The possible range is reported to be 2.5°C to 7.8°C (IPCC 2014a: 8).
5 These and other risks are detailed in IPCC 2014c, especially Part B ("Future Risks and Opportunities for Adaptation").
6 Estimates of the cost of mitigating climate change depend partly on how much warming we are willing to accept. According to the IPCC (2018), limiting warming by 2100 to 1.5°C is apt to be considerably more costly than limiting warming to 2°C; on the other hand, adapting to 1.5°C is apt to considerably less costly than adapting to 2°C. Much controversy surrounds attempts to calculate costs and benefits of climate change; see Broome (2012), especially Chapters 3, 8, and 9, and the references given therein.
7 See, for instance, most of the papers in Moore and Nelson (2010). In their introduction to the volume, Moore and Nelson write "[w]e have a moral obligation to avert harms to the future, so as to leave a world as rich in life and possibility as the world we inherited."
8 Arguably, there are some future people, especially in the near future, who would exist whether we mitigate climate or not, for the non-identity effect is time-lagged (Parfit 1976: 102). However, those people are unlikely to benefit much from mitigation, because the effects thereof (in terms of avoided warming) are also time-lagged (USGCRP 2017: 394). Tebaldi and Friedlingstein (2013) find that it would take 25–30 years for the effects of mitigation to become discernible. Likewise, the IPCC (2014b: 9) projects a similar increase in global temperature across all emission scenarios over the next few decades. It is plausible, therefore, that the time-lags of non-identity and mitigation would approximately offset. I thank an anonymous referee for raising the issue.
9 For instance, the World Health Organization estimates that global warming would contribute an additional 250,000 deaths annually by 2030 (Hales et al. 2014).
10 The example is borrowed from Northcott (2008), who adapted it from Lewontin (1974).
11 One interesting question we might ask is whether Twin Earths in the same "treatment group" (i.e. the same cell in the table, subjected to the same settings of the various exogenous variables) wind up with the "same" individuals (i.e. twins) being born? To put it another way, how much of the variance is due to random chance, or factors for which we cannot control? Perhaps quite a bit.
12 More precisely, these are the requirements to keep the temperature increase below 1.5°C "without overshoot". IPCC (2018) also explores pathways where we temporarily exceed 1.5°C above preindustrial average and then draw down by means which are technically possible but currently unavailable.
13 Thirsty Traveller is adapted from Mackie (1974: 44), who offers it as a counterexample to a simple counterfactual account of causation.

References

Ariew, André (1996): Innateness and Canalization. In: Philosophy of Science, 63 (5), 19-27.

Bontly, Thomas (2016): Causes, Contrasts, and the Non-identity Problem. In: Philosophical Studies, 173 (5), 1233-1251.

Broome, John (2012): Climate Matters: Ethics in a Warming World. New York: W. W. Norton & Company.

Caney, Simon (2005): Cosmopolitan Justice, Responsibility, and Global Climate Change. In: Leiden Journal of International Law, 18 (4),747-775.

Gardiner, Stephen (2006): A Perfect Moral Storm: Climate Change, Intergenerational Ethics and the Problem of Moral Corruption. In: Environmental Values, 15 (3), 397-413.

IPCC (2018): Summary for Policymakers. In: Global Warming of 1.5°C. Edited by Masson-Delmotte, V. et al. https://www.ipcc.ch/site/assets/uploads/sites/2/2019/06/SR15_Full_Report_Low_Res.pdf. Viewed 15 December 2019.

IPCC (2014a): Summary for Policymakers. In: Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Edited by: Edenhofer, O. et al. Cambridge: Cambridge University Press

IPCC (2014b): Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Edited by Pachauri, R.K. / Meyer, L. A. Geneva: IPCC.

IPCC (2014c): Summary for Policymakers. In: Climate Change 2014: Impacts, Adaptation, and Vulnerability. Part A: Global and Sectoral Aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Edited by Field, C.B. et al. Cambridge: Cambridge University Press.

Lewontin, Richard (1974): The Analysis of Variance and the Analysis of Causes. In: American Journal of Human Genetics, 26(3), 400–411.

Mackie, J. L. (1974): The Cement of the Universe. Oxford: Oxford University Press.

Mill, John Stuart (1858): A System of Logic. New York: Harper & Brothers.

Moellendorf, Darrel (2015): Climate Change Justice. In: Philosophy Compass, 10 (3), 173-186.

Moore, Kathleen / Nelson, Michael (2010): Moral Ground: Ethical Action for a Planet in Peril. San Antonio, TX: Trinity University Press.

Northcott, Robert (2008): Can ANOVA Measure Causal Strength? In: The Quarterly Review of Biology, 83 (1), 47-55.

Parfit, Derek (1976): On Doing the Best for Our Children. In: Bayles, Michael D. (ed.): Ethics and Population. Cambridge, MA: Schenkman Publishing Company Inc., 100–115.

Parfit, Derek (1984): Reasons and Persons. Oxford: Oxford University Press.

Putnam, Hilary (1975): The Meaning of "Meaning". In: Minnesota Studies in the Philosophy of Science, 7, 131-193.

Roberts, Melinda (2003): Can It Ever be Better Never to Have Existed at All? Person-Based Consequentialism and a New Repugnant Conclusion. In: Journal of Applied Philosophy, 20 (2), 159-185.

Sober, Elliott (1988): Apportioning Causal Responsibility. In: The Journal of Philosophy, 85 (6), 303-318.

Tebaldi, Claudia / Friedlingstein, Pierre (2013): Delayed Detection of Climate Mitigation Benefits Due to Climate Inertia and Variability. Proceedings of the National Academy of Sciences, 110 (43), 17229-34.

Tremmel, Jörg (2018): Fact-insensitive thought experiments in climate ethics – Exemplified by Parfit's non-identity problem. In: Jafry, Tahseen (ed.): The Routledge Handbook of Climate Justice. London: Routledge, 42-56.

USGCRP (2017): Climate Science Special Report: Fourth National Climate Assessment, Vol. 1. Edited by Wuebbles, D., et al. Washington, DC: US Global Change Research Program.

WHO (2014): Quantitative risk assessment of the effects of climate change on selected causes of death, 2030s and 2050s. Edited by Hales, S. et al. Geneva: 2014.

WHO (2009): Global health risks: mortality and burden of disease attributable to selected major risks. Edited by Mathers, C. / Stevens, G. / Mascarenhas, M. Geneva: 2009.

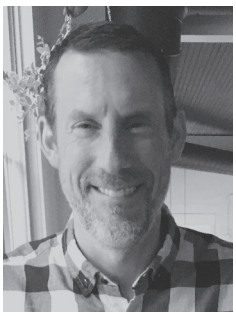

Thomas D. Bontly is an associate professor of philosophy at the University of Connecticut. He has research interests in the philosophy of mind, philosophy of science, ethics, and environmental ethics. His current research focuses on the non-identity problem.

Contact details:
Thomas D. Bontly
Philosophy Department, U-1054
University of Connecticut
Storrs, CT 06515
USA
E-Mail: thomas.bontly@uconn.edu