The ethical issues concerning climate change are very often framed in terms of harm: so people say that our acts (and omissions) affect the environment in ways that will cause *severe harm* to future generations, or that we are using natural resources at the *expense* of future generations, making them *worse off*. Quite a lot of people in this room have pointed out that part of the difficultly with addressing climate change and carbon emissions is that if they cause harm, they cause it in a manner that's very different from paradigmatic ways of harming others. So we heard this morning again about discussion of causal effects and problem about agency since if harm is caused because it looks as if it's a result of the combined actions of very many, million peoples acting together via unprecedented technological pausing. There's a lot more that could be said about this important problem.

I'm going to talk about a more radical way in which the ethics of climate change seems to depart from commonsense morality and that's because it not only concerns acts that effect people who don't yet exist—i.e. future generations—but also acts that determine *who* comes to exist. And this generates what the oxford philosopher Derek Parfit famously called a non-identity problem. The problem raises, I think, a more fundamental challenge to the common assumption that what we do now or don't do now can really be said to harm future generations.

So here is what I'll try to do. I'll introduce the non-identity problem and I know a lot of this will be extremely familiar to some of you and far less so to others, so I'll be quick but I hope not too quick. I'll first first introduce the problem the context of reproductive ethics where it's received the most attention. Well I've done quite a bit of work on this issue with Julian Savulescu. Then I'll explain how it also arises in a slightly different way in the context of climate change and I'll review some out of very many solutions that have been proposed to this problem, though not remotely an exhaustive survey. But the first part is really intended to introduce you to this problem and give you sense of its practical importance but also a sense of why after a philosophical discussion of over around 40 years there is still no agreed solution to the problem. People keep coming with new approaches and solutions.

I hope I'll have time at the end of the talk to also say something more original and substantive about the issue of moral uncertainty. The non-identity problem is one example of what John Broome in his book *Climate Matters* calls the moral certainty arising from a lot of ethical question surrounding climate change. I'll briefly say something about how Broome proposes to address this problem, but I'll end with a more specific proposal relating to the non-identity problem. For many years I've been thinking and working about the non-identity problem, as I said above, in the reproductive context, and I've always thought that if there is one important simple, significant philosophical example that should be more widely known it is the non-identity problem. But preparing this talk led me to thinking that that's actually a mistake and I'll end by some remarks suggesting it might be better for philosophers to keep quiet about the non-identity problem. I'll explain later what I mean by that.

So the non-identity problem arises from this point about our existence that some people call it precariousness of existence, the fact that who actually comes to exist at any point in time is a highly contingent matter, a matter of great luck. Even tiny adjustment to the process leading to the conception of each of us would have led to someone else coming to exist.

And it's a very doubtful to say the least that we would still come to exist if some other sperm fertilized another ovum or if our parents never met or their grandparents never met or there was no industrial revolution. And this straightforward but remarkable empirical fact generates some surprising conclusions. You can see that if you think about this example from a reproductive ethics. Imagine there is a rubella epidemic and if a couple conceives the child now the woman will contract rubella and the baby will be born with congenital rubella, blind and deaf, so severely disabled. But in a few months, the epidemic will pass and the couple could have a healthy child. And it seems deeply wrong for the couple not to wait. And the problem is to explain why because not waiting to have a child later will lead to an outcome of a child with severe disability but it also determines who comes to exist. So if the couple waits longer, they will have a healthy child but that of course will not benefit in any way the disabled child they would have if they don't wait because it would lead to the existence of *another* child. So that disabled child that they conceive have now will be in no way benefited that or made better off by the couple waiting longer. And it sounds like we can't really say that this child is harmed by the couple's decision not to wait longer.

And we can set it out as a kind of formal argument. It starts from the seemingly plausible assumption that an act is wrong only if it harms someone, and given that if the couple refuses to wait, this seems to harm no one because it doesn't harm the resulting disabled child since it wouldn't exist otherwise, and it can't harm the alternative healthy child they could have had because it doesn't even exist.

And these premises seems to lead us to a surprising conclusion that couple's action is not wrong in anyway, a conclusion that many people find hard to accept. In the productive context there are many even more direct cases where we can choose which child we're going to create and the same problem arises.

Again, the reproductive context has received a lot more attention than the context of the environment and climate change to which I now turn.

So just getting back to this point about the precariousness of existence. It also applies in the context of large scale policy or collective activity, for example to mitigate climate change. Because, as Derek Parfit already noted many decades ago when he first introduced the problem in a systematic way, this kind of large scale policy affects the identities of who will later come to exist almost right away because they will affect everyday life, changing which couples meet and marry and if and when they procreate. So over time completely different people come to exist simply because we adopted that policy.

This means that many of the people who will suffer if we don't adopt such a policy if we do nothing about climate change wouldn't even come to exist if we do adopt it. Again, we have two choices, and if we do nothing in this case, this will lead to a very bad outcome in one sense—all sorts of suffering and death and hardship—but this very act also determines who comes to exist. So if we make great sacrifices to prevent climate change or mitigate it, the result will be better but it will not be better for the people who would have existed if we had done nothing, and it looks like if we do nothing it doesn't look like we actually harm them in any straight forward sense.

Again you can spell this out as an explicit argument, which essentially runs exactly like the case of the disabled child in the rubella example. The climate change case is slightly different because it's not a pure case because some of the harm we cause now and which we can prevent actually applies to people who already exist, people who are be alive in 20 years or 30 year or just about to be born are likely to also suffer the time, so the conclusion is not going to be that by doing nothing we are doing nothing wrong but it does seem to suggest that by not making great sacrifices we will be acting in a way that's far less wrong than we may think.

Now there's the interesting empirical question of how quickly identities will be affected by different policies that so far as I know hasn't been seriously addressed and that's bit of puzzling given the importance of the issue. It's important to bear in mind that even if many things we do or don't do will harm certain people in 20 or 30, 40 years, quite a lot of the people who will exist at that point in time wouldn't exist if you don't adopt these policies, so we are harming less people than we may think. The problem doesn't only kick in, in let's say 100 or 150 years when everybody has been 'replaced'.

Another point to point out that's quite different from reproductive cases is that the different policies we adopt arguably would have affect identities over time more quickly or more slowly. So of course if it would involve everyone eating a lot less meat over the world this would pretty much immediately effect who comes to exist next but if we adopt policies like geo-engineering, arguably that effect on identity would take quite longer to kick in and that this may have interesting ethical implications, but I don't have time to discuss that. I just wanted to point out that it's odd that these issues haven't yet received much attention.

Ok. So the challenge presented by the non-identity problem is to explain how it could be wrong to do nothing to prevent that adverse effect of climate change if, as this form of argument suggests, doing nothing couldn't be said to harm future generations. And this isn't a conclusion that many people are keen to endorse.

Just quickly, let me mention that there is also a backward application of this problem in the context of the environment and climate change. It is very common that a lot of people in the west have greatly benefited from the industrial revolution and its aftermath but that these past events have also had a great effect on the environments and have harmed and will harm people in developing countries and that that's why we have a duty to remedy that harm. People often call this the beneficiary pays principle.

But again if you think about this consideration relating to identity, it seems very doubtful that we in the west could be said to have benefited from the industrial evolution and its aftermath because if it hadn't happened, we wouldn't even exist, and similarly for the people supposedly being harmed or who supposedly will be harmed. So the problem applies not only in the direction of the future but also backwards.

So it's kind of very simple philosophical problem, kind of philosophical discovery you might even think that it's pretty easy to set out and has all kind of surprising and even disturbing practical implications. Very many people have tried to solve or address the problem. I'm not going to review all of these solutions, some of which are extremely complicated and difficult to explain, but I'll quickly review just a few of them just to show you how different ways of dealing with the problem will have quite different practical implication in the context of climate change.

One approach—it's not terribly popular—is called the person-affecting view, essentially follows the slogan that Narveson first put forward some years ago, that morality is about making people happy and not about making happy people. On this view, morality is person-affecting, that is to say, the only thing that only matters is whether we harm or benefit actual people (including future people who will exist anyway). So to bring to existence this person rather than that person as long as we cannot be said to harm or benefit someone is a morally neutral decision.

Essentially this kind of view bites the bullet, if you want, and just endorses the conclusion of that surprising argument, and people who endorse this view essentially conclude, however reluctantly, that it's actually not wrong to deliberately create a disabled or unhealthy child when we can just as easily create a healthy one.

Not many people endorse this view. In the context of the environment, it would have pretty dramatic practical implications, suggesting something like the counterintuitive conclusion that we have seen earlier that we should make far fewer sacrifices to prevent long-term climate change and focus instead on addressing shorter term effects that apply to the people exist today rather than on people existing in the future.

Like I said, very few people willing to go that way. It's perceived as profoundly counterintuitive. Although a lot of people tend to be quite sceptical about all kinds of moral intuitions, I think this particular intuition—that such acts must be wrong even if they don't seem to harm anyone—hasn't really received very serious critical scrutiny and I'm not sure it would withstand that scrutiny. I myself find this conclusion extremely hard to accept, so I'm applying this last point to myself as well.

Another view is total utilitarianism, which tells us that we should maximize future goodness, whether or not we benefit or harm anyone. And the fact that our acts also affects who comes to exist is on this view morally irrelevant, so this is essentially is very impersonal view of ethics. It doesn't really matter which particular person you benefit or harm, you should just try to make the world better in something impersonal sense.

If you hold this view, you certainly don't need to worry about the non-identity problem. This view does imply that we should make very demanding sacrifice as now to make the world impersonally better in the distant future even if that won't benefit anyone.

Like I said, this view would get us out of the counterintuitive implications of the non-identity problem, but it generates many other problems of its own, utilitarianism generally is not a terribly intuitive view. There are also specific problems that come up if you hold that view in the environmental context or reproductive context: total utilitarianism seems to imply that we have an obligation to create more children. So even if a healthy child might be impersonally better than a disabled one, a disabled one would still be impersonally better than no child at all. And it seems quite contrary to commonsense to think that we have an obligation to create more and more children including disabled ones. Total utilitarianism also generates what Parfit famously called the repugnant conclusion, very briefly, the implication that it would be better if you created numerous lives barely worth living instead of a much smaller number of very good lives.

Parfit himself suggested a kind of provisional solution. Yes, we should choose the impersonally better outcome but only in contexts where we are just choosing between same numbers. This just sidesteps the problem of procreation itself, as well as the repugnant conclusion. I won't get into the details but it's obviously an arbitrary restriction and as he himself is very explicit about it's just a provisional response to the issue.

But anyway in the context of the environment this solution doesn't really help. In reproductive ethics we often can choose between creating this child or that child but in the environmental policy context whatever we choose to do (or not to do) will inevitably also affect the numbers of people who come to exist, and the problems of population ethics can't really be avoided.

The final view I'll mention and I think it's getting slightly more attractive or more popular—I'm attracted to it myself—is that we should try to combine elements from both of those approaches. So on that view it's more important not to harm (and to benefit) people who exist (or who will exist anyway). But we also have *some* reason to promote impersonal goodness or prevent impersonal badness. Applied to the context of climate change, the view seems to suggest (but this depends on how you spell it out) that we have reason to make some sacrifices to make the world better in the distant future but we should focus more on present suffering and short-term effects. One problem, however, is that if there is tons of impersonal value in the long-term future this would swamp out any person-affecting concerns that relate to existing people or people who will come to exist fairly soon. So this 'hybrid' view may ultimately collapse into something indistinguishable from the total utilitarian view with all of its problems.

Anyway, these just a three out of very many attempts to respond to the problem. There are many other views that I will not have time to discuss, but each of them has serious problems. And some of them anyway only apply in the reproductive context and would not make sense or have any application in the context of climate change.

So I've tried to give you a sense of what the problem is, why it's practically important and also of the general state of philosophical discussion. So as I said, after around 40 years of discussion there is no consensus as to how to deal with it and what the practical implications are going to be.

Now I had all sorts of interesting things to say about how John Broome spells out the non-identity problem and its implications in his interesting and very good recent book *Climate Matters* but I'll turn, in the few minutes I still have, to the more significant point I want to make in this talk. It is not philosophical but more sociological but it has some practical significance.

So one thing that's a bit disappointing about John Broome's book is that he makes very many interesting philosophical points and arguments but then highlights the fact that just as there is significant empirical uncertainty in the context of climate change, there is also a significant degree of moral uncertainty because a lot of unresolved philosophical and ethical questions will very dramatically affect what we ought to do in the environmental context. And the non-identity problem is just one instance of such uncertainty that Broome mentions; there are quite a few others relating to population ethics, questions about value and how to aggregate it, the value of life etcetera, etcetera.

What follows from that? Well, one thing that Broome emphasizes and that seems like a very plausible and attractive point to many people in the room, I think, is that policy needs a lot more input from moral philosophy. So as he says "we shall be guided by the quantitative analysis of our economists and scientists, but we must make sure that our analysis rests on good moral foundations". There are a lot of normative assumptions and positions that need to be make explicit and scrutinised. So there is an important role for moral philosophy here.

But because there is also significant moral uncertainty and disagreement within moral philosophy, moral philosophers and ethicists cannot offer authoritative prescriptions. So Broome suggests that what philosophers should do instead is just introduce policymakers and the general public to the main ethical issues, explain the different moral considerations and arguments, and make the general public and policymakers aware of the different philosophical options. And I think I've heard over the years, Julian Savulescu present this as a general attractive view of what practical ethics is about: not telling people what they ought to do but introducing them to the ethical and philosophical complexity that is often underestimated by ordinary people and policymakers.

Another thing that John Broome suggests is that because of this moral uncertainty and because there is no normative consensus about all these difficult philosophical issues, agreeing about what we ought to do about, say climate change, it's critical that it's done through a democratic procedure. As he says, "democratic public debate is a present the only means we have of coping with this sort of moral uncertainty."

Now there is a lot more to say about this. I think quite a few people in the room will probably raise eyebrows and quite a few people here have pointed out that democracy is not really the solution but very much part of the problem here. Again we've heard about this earlier this morning and I won't go into that in detail but I want to focus more specifically about the question of how we ought to deal with the moral uncertainty involved with the non-identity problem. Broome's idea is that we should explain the non-identity problem and its implication and the different approaches and options on the table to the wider public and to policymakers is you might think is clearly an important task for moral philosophy.

And as I said, this is something I also used to think until I prepared this talk. And it is a very striking fact that although the non-identity problem has been discussed by philosophers for, I would say 40 years—I'm not sure exactly but it's around 40 years—in lots of articles and books. It applies in important ways to productive context and to climate change and environment, yet it's still barely known outside of philosophy. It's a striking and maybe an embarrassing fact for moral philosophy.

I've tried to Google 'climate change and non-identity' and pretty much everything that comes up are philosophical articles and books, including by people sitting here, but there is almost nothing beyond that apart from one social science article that does mention the problem but completely misunderstands it. So it's very widely discussed within philosophy and almost entirely unknown outside even in the reproductive context where it's quite easy to explain and its implications are very straightforward. You see law and policies routinely phrased in ways that make no sense given the non-identity problem—in the UK, for example, it's wrong to select for this because it's not in the interest of the child.

Kahane

So you might think that what we need to do now is to introduce the non-identity problem to the general public, describe it, and explain the different options—essentially what John Broome does in his book. But I have a worry about this. While the most progressives agree that significant sacrifices need to be made to tackle climate change, as you know, very many others resist this conclusion.

Now so far much of this resistance has appealed to supposed (and often illusory) empirical uncertainty, essentially ignoring the scientific consensus. I think that kind of strategy would be hard to maintain over time and worries about moral uncertainty—about normative and political issues—are going to become far more important to debate about climate change than they have so far.

And it seems to me like plausible sociological speculation that the non-identity problem, if known widely, would be a pretty powerful tool or powerful 'rationaliser' in the hands of many people who want to do little or nothing about climate change. After all, the non-identity problem does appear to show that we can't harm future generations if we don't make great sacrifices, or even if we do nothing. And within philosophy there is no consensus as to how acts that lead to bad outcomes without harming people could be wrong. And although the person-affective view I discussed earlier, as I said, is not particularly popular with the largely liberal moral philosophy, I would predict that it will be far more popular among those who wish to avoid making significant sacrifices once this issue becomes more familiar to people outside ethics.

Of course many philosophers think that they do have a solution to the problem but such solutions will always take the form of something like 'what we do can't really harm later generations, but it's still wrong because of [insert some weird, complex and controversial philosophical theory]'. And the first part of this story is extremely easy to grasp and accept—it doesn't require any very sophisticated or elaborate or controversial philosophical background to explain the non-identity problem. But the second part of the philosophical story isn't going to be so straightforward. Any theory that proposes a solution to the non-identity problem is going to be a lot more complicated controversial and far easier to dismiss. And of course there is no empirical data, let alone knock down philosophical argument, that could refutes the person-affecting view in a way that would be widely accepted.

Of course what I just said only applies to philosophers who think we *should* be willing to make significant sacrifice to reduce the adverse affects of climate change, and that future generations do matter. And the tentative conclusion of this talk is that these philosophers should really keep quiet about the non-identity problem! So we shouldn't write popular books explaining the problem or mention it in policy papers or public discussion not to mention even more direct ways of engaging the general public, like TED talks, blogs or the popular media… And I'll stop here. Thank you.