

A reply to my critics

Midwest Studies in Philosophy, 40 (2016), pp. 158–171

1 Introduction

Climate Matters is aimed at a wide readership. When I was commissioned to write the book, I thought that people would be interested in what moral philosophy has to say about the private morality of climate change. How should we as individuals respond to the existence of climate change? I had not previously thought much about this question, but I drew my own conclusions about it and presented them. I surprised myself by the conclusions I came to.

I never thought the world should try to deal with climate change by promoting private morality. That would be hopeless. Far too few of us will respond as we morally should, and those who do will have little effect. An effective response to climate change will have to come from governments, who can use their powers of tax and regulation to influence the behaviour of very large numbers of people. So *Climate Matters* was principally aimed at public morality: how governments should respond to climate change.

Since our governments must act, this gives us a derivative moral duty to try and get them to act rightly. I call this ‘civic morality’ and distinguished it from private morality; it is the morality that is asked of us in our role as citizens.¹ What we can achieve as citizens may be much more important than what we can achieve as private individuals. *Climate Matters* principally aims to put citizens in a position to understand the moral principles that underlie what governments should do.

I should have known that, when moral philosophers came to read the book, many of them would be most interested in what it says about private morality, since private morality is their main business. And I should have known that, when political philosophers came to read it, many of them too would be most interested in the same thing, since their main business is justice and I argued that the private morality of climate change is principally regulated by justice. It is also true that my conclusions about private morality surprise other people besides me, and that naturally draws their attention. Two of my commentators in this volume have concentrated on this topic. I am very pleased and grateful that philosophers have troubled to read my book at all, and I especially appreciate the work of those who gone so far as to write responses. We had an excellent conference at Essex, organized by Paul Bou-Habib, to discuss the comments that appear in this volume. I appreciate the kind remarks the authors have made about my book. So this is very far from a complaint. But I want to take this opportunity to reiterate that what really matters in climate change is what our governments do about it. We should do what we can to influence them for the better.

2 Cripps

Still, I start with private morality. My view about this is that none of us should emit any greenhouse gas in total. We may achieve this aim by reducing our emissions and offsetting any emissions we still make. The reason we should not emit is that justice requires us not to harm other people, and emissions of greenhouse gas harm other people.

This conclusion arises from the empirical facts of climate change. I am going to stress two of these facts, but before I come to those, I need to mention a sort of proto-fact. The harm done by emitting greenhouse gas is done only through the effect it has on the global concentration of greenhouse gas in the atmosphere. This is true, not only of the harm done through climate change, but also of the harm done through the acidification of the oceans. The harm done is determined through this one quantity – the global concentration of greenhouse gas – and its development over time.

This means that you do no harm by your emissions if you do not add to the global concentration. If you emit at one place, and also prevent an equal quantity of emissions at another place, you do no harm because you do not change the global concentration. This is how offsetting works. Elizabeth Cripps asks whether harm will ‘be prevented to exactly the same people as our individual emissions would have harmed’. The answer is ‘Yes, of course’. If you do no harm, you harm no one. Offsetting is not a matter of compensating for harm done by reducing other harm. It is a matter of not doing harm in the first place.

The proto-fact is slightly approximate. There are several different greenhouse gases, and they behave differently in the atmosphere. It is only an approximation to lump them together into a ‘carbon-dioxide equivalent’ concentration, as we commonly do. The proto-fact applies more exactly (though still not quite exactly) to carbon dioxide than to other greenhouse gases. Carbon dioxide stays in the atmosphere for a very long time and is very well mixed around the globe, so its effect does not depend on where it is emitted. Offsetting is therefore more reliable for carbon dioxide than for other gases.

Now the first of my two empirical facts: the harm done by greenhouse gas emissions is proportional to the quantity of emissions. For example, a billion tonnes of gas do a thousand times as much harm as a million tonnes.

This remark needs three qualifications. First, on a very large scale the harm done is thought to be more than proportional to the quantity of emissions. A trillion tonnes of greenhouse do more than a thousand times as much harm as a billion tonnes.

Second, on the much smaller scale of, say, thousands of tonnes, the harm is lumpy. Some of the harm done by climate change results from discrete events. A baby dies of cholera, or a flood destroys a city. So if you drew a graph of the harm done against cumulative emissions (the total of greenhouse gas that has been emitted), it would show upward jumps. The position of these jumps is unknown and, because the atmosphere is a chaotic system, unknowable. When I say the relation between harm and emissions is proportional, I am speaking of the graph after smoothing out the lumps. More exactly, I am speaking of expected harm rather than actual harm.

Third, even on a large scale, the relation between harm and emissions is uncertain. For one thing, there could be an upward jump on a very large scale. There could be some tipping point for runaway climate change so that, if cumulative emissions go beyond some particular point, climate change will accelerate uncontrollably and vast harm will result. If there is a tipping point like this, its position is uncertain. For another thing, our best science leaves the amount of harm done per tonne of gas uncertain. I shall discuss what this amount is in section 3, when I turn to my second empirical fact about climate change. So even on a large scale, I am strictly talking about expected harms rather than actual harms.

To summarize, the empirical fact, stated more precisely, is that, on all but the largest scale, the expected harm done by emissions is proportional to the quantity of emissions. Emissions do expected harm at a constant rate per tonne. One tonne of emissions does one billionth of the expected harm done by a billion tonnes.

The expected harm caused by any group of individuals is the total of the expected harm caused by each of them separately. Elizabeth Cripps (section III) ‘flags up’ ‘the distinction between a number of individuals performing actions harmful in themselves which add up to greater harm, and a number of individuals performing actions which make no difference at all in isolation but, in combination, bring about serious harm’. She does not say clearly that climate change is on the former side of this distinction, but that is the fact.

At least, it is the fact if we read ‘harm’ as ‘expected harm’. What about actual harm? Could it be true that an individual’s emissions do no actual harm? No. Some of the harms done by

climate change come in discrete lumps, but others are continuous. For example the gradual rise in sea levels caused by climate change will steadily erode the land, and in other ways farming in many areas of the world will progressively become more difficult. Many people will have to work harder and will find themselves hungrier. Many will have to walk further to fetch water. Climate change will make the environment progressively less pleasant in many parts of the world. Summer weather will be hotter, for example, and nature will be less beautiful. Any increase in greenhouse gas will cause an increase in these continuous bad effects even if it does not precipitate one of the lumpy harms. So each person's emissions will lead to actual harm as well as expected harm.

What does the first empirical fact tell us about the private morality of climate change? Suppose, first, that you aim to promote good in the world. When there is uncertainty, so you do not know exactly how much actual good will result from your acts, the theory of value tells us you should promote expected good instead, and avoid expected harm. By reducing your emissions you can reduce expected harm at whatever is the constant rate per tonne. If you aim to promote good, this is a way of doing it. It gives you a simple moral reason to reduce your emissions.

It means the morality of climate change is simpler than some philosophers have thought. These philosophers assume that climate change is on the latter side of Cripps's distinction. They assume that each individual does no harm by her emissions.² Since they nevertheless think we ought to reduce our emissions, they have to find a reason for reducing them even though they are harmless. This is difficult. It is an unnecessary difficulty they have brought on themselves.

However, I argued in *Climate Matters* that the reason I just described for reducing emissions – that it promotes good – is not a sufficient reason for doing it.³ There are better ways of using your resources to promote good. One is to support a charity that treats tuberculosis, for example. I argued that the reason why you should reduce your emissions is not that one. Instead, it is a reason of justice. Justice requires you not to harm other people, at least not for your own benefit. Since emissions of greenhouse gas do harm, you should not make them.

You might doubt that this principle of justice applies to expected harms as well as to actual harms. I explained that your emissions will cause actual harm as well as expected harm, and this is enough to make them an injustice. But since the actual harm may be only a small part of the expected harm, it is worth considering what justice has to say about the imposition of expected harm too. Does justice prohibit you from doing expected harm as well as actual harm? Since imposing expected harm is imposing a risk of harm, the question is whether justice prohibits you from imposing a risk of harm.

Suppose you drive dangerously down a street, imposing a risk of harm on pedestrians, but luckily you hurt no one. Evidently you do something wrong. Moreover, it seems intuitively that the wrongness is not merely that you diminish the expected value of the world. It seems you wrong the pedestrians particularly, by putting them in danger. So it seems this wrong is an injustice done them. How can we make sense of that?

In any of several ways. One is to say that a risk of harm is a sort of actual harm. So in this case you actually harm the pedestrians. If that is true, then the principle of justice that you should not harm other people applies directly to expected harm. A second way is to say that, as well as the duty of justice not to harm people, there is another duty not to impose a risk of harm on them. A third is to say that you ought not to take a risk of doing someone an injustice, even if you luckily do not actually do her an injustice. This would be a justice-related wrong done a person, though not an actual injustice.

The first of these options seems to me plainly false. In *Climate Matters* I declined to choose between the second and third of them,⁴ but I now lean towards the second. The third leads to the same conclusions as the second, but by a more complicated route. So for simplicity I shall adopt the second. I assume that emitting greenhouse gas is an injustice done to other people because it exposes them to a risk of harm, and that this would be so even if you did them no actual harm.

I conclude it is unjust to emit greenhouse gas because of both the actual harm and the expected harm it does to other people. The injustice you do each person is small because the actual and expected harm you do her is small. But because you harm a great many people, it adds up to a serious injustice.

A number of philosophers argue that injustices do not add up across people.⁵ But if I understand her, this is not Cripps's objection to my view. Instead, she denies that the harm you do each person is 'significant'. But that term begs the question. A harm is significant when it matters. You cannot assume a quantity is insignificant just because it is small. Is a penny significant? Not if you lose it by accident. But if a bank steals a penny from each of its millions of customers, it is very significant. The bank perpetrates a serious injustice. It does a small injustice to each person, and these small injustices add up to a big one. In the case of climate change, the harms you do to other people matter even though they are small, because there are so many of them and they add up. So they are significant.

Philosophers sometimes seem to think of a very small number as much the same as zero. But there is an extremely significant difference between a very small number and zero. If you add a great many zeros you get zero, whereas if you add a great many small numbers you get a big number.

Cripps also makes the different point that the harm you do each person by your emissions is imperceptible. She is right, unless your emissions happen to precipitate a discrete harm such as someone's early death. But Derek Parfit explained long ago that a harm does not stop being a harm just because it is imperceptible,⁶ and Cripps does not claim it does. Moreover, imperceptible harms may be significant. So perceptibility is irrelevant.

Cripps finally points out that the injustice we each do by our emissions (if I am right that there is one) has a very different character from the injustices we typically think of in philosophy. We typically think of large harms done to a single person rather than small harms done to many people. So we should not take for granted the conclusions we commonly draw about injustice. In particular, she says, we should be cautious about claiming that the duty of justice not to emit greenhouse gas 'takes priority over any duties of goodness we have, such as our duty to promote government action on climate change'. I agree. Cripps attributes this claim to me, but I cannot find where I made it. In my chapter on private morality, I particularly excluded any discussion of our civic duty to promote government action on climate change.

As a brief digression before coming to my second empirical effect, I want to mention a different and very interesting issue that Cripps raises. It is an issue of attribution. Justice requires you to offset emissions that are correctly attributable to you, but correct attribution is often difficult and controversial. It is controversial not only philosophically but politically. Consumers in western countries buy goods produced by manufacturers in eastern countries. Greenhouse gas is consequently emitted. Both the consumers and the producers cause the emissions; the consumers are a more distal cause and the producers a more proximal one, but both are causes. Should the emissions be attributed to the western countries or the eastern ones? This is a political issue.

I am sorry to say I have no formula for determining attribution. But so far as justice is

concerned, I think we must recognize one constraint on it. The total of emissions attributed to each person should be equal to the total actually emitted. For example, when greenhouse gas is emitted in the course of producing a consumer good, we should not attribute this emission to both the consumer and the producer. The emission should be offset only once.

Cripps raises the issue of attribution between parents and children. When a child emits greenhouse gas, both she and her parents cause the emission. The child causes it more proximally; the parents cause it more distally by causing the child to exist. Should we attribute the emission to the child or the parent? In this case, it seems to me fairly clear that we should attribute it to the child, once she becomes an autonomous person. That seems to be an implication of taking her to be autonomous. We therefore should not attribute the emission to the parents. Justice requires the child to offset her emissions, not the parents.

Any duty of justice that falls on the parents is more remote. Perhaps they should teach the child to act justly. Perhaps they should make sure the child is in a position to perform her duty not to emit.

3 Lawford-Smith

Now my second empirical fact about climate change. This one is quantitative. Holly Lawford-Smith makes the case that quantities are important because we have to balance different considerations against each other. I agree entirely.

The second fact is that the cost of preventing greenhouse gas from getting into the atmosphere is far less than the amount of harm it does once it gets there. Figures for the cost and the harm can only be rough, but rough figures are enough to demonstrate this fact. The cost to people like you and me of preventing emissions is in the region of \$10 per tonne of carbon dioxide. This is about the price that reputable companies charge members of the public for offsetting our emissions. I assume they do their job.

The harm done by emissions is much harder to put a figure on. Climate change is harmful in multifarious ways. For instance, it impoverishes the natural world and it is very damaging to particular human cultures. I am thinking especially of the cultures of the native people of the Arctic, the forests and small islands. These harms and others cannot plausibly be measured in monetary terms. But the harms of climate change also include ones that economists have set a value on. They call them 'the social cost of carbon', and have estimated how much this is.

Estimates of the social cost of carbon depend strongly on the discount rate that is assumed. Since many of the harms occur far in the future, a lower discount rate leads to a higher estimate. For less obvious reasons that come from the workings of the models used in making the estimates, they also depend on the date when the gas is emitted. Gas emitted later turns out to have a higher social cost. For emissions in 2015, the US Government estimates the social cost of carbon, per tonne of carbon dioxide, at \$42 for a discount rate of 3% and \$65 for a discount rate of 2.5%.⁷

These figures must be regarded as very serious underestimates of the harm done by carbon dioxide. They basically aim to measure the reduction in the world's economic output that will be caused by climate change. They do not include damage to nature, animals and human cultures. They do not even include the direct harm people will suffer as a result of living in an impoverished environment. They do not include one of the most serious harms climate change will do to people, which is to shorten people's lives through disease, poverty, natural disasters and in other ways. In *Climate Matters* I mentioned this harm of shortening lives in order to give readers an idea of the amount of harm their emissions do.⁸ Each of us by our lifetime emissions will shorten people's lives by a few months in aggregate. This is bad

enough on its own to make it the case that we should not emit, but it is not even included in the US Government's estimates of the social cost of carbon.

Furthermore, since the estimates are based on aggregate economic models, they do not reflect the fact that the harm of climate change will fall more heavily on the world's poor than on the rich. Since money matters more to the poor than to the rich, this too means that harms are undervalued. It is also arguable that even 2.5% is too high a discount rate to use.⁹ All in all, a more accurate figure for the social cost of carbon might be several times greater than the ones I have just given.

This means that, for a cost of \$10 spent on offsetting your emissions, you can save yourself from doing a much, much greater harm. Even the economically measured part of the harm you avoid is five or more times greater. Given that, I think it is obvious that you should do it.

One response to this conclusion of mine is to deny that offsetting is truly a way to avoid doing harm. But that is not Lawford-Smith's response. Instead, she points out that you cannot avoid doing some harm in your life. No doubt she is right. But why should this affect my conclusion about offsetting emissions, which is something you *can* do to avoid doing harm? Lawford-Smith says it shows that you have to balance different ways of avoiding harm against each other. I argued in *Climate Matters* that you should offset your emissions simply because harming other people is an injustice and you therefore should not do it. But she thinks the idea that you should not do harm is too simplistic because you cannot avoid doing harm. We have to replace the simplistic aim of not doing harm with the more sophisticated aim of minimizing harm.

Does this really affect my conclusion about climate change? You have the opportunity, at a small cost to yourself, to avoid harming other people to a very, very much greater extent. Could it conceivably be morally permissible for you not to lay out \$10 on offsetting, when this would save you from doing \$50 worth of harm to others, and probably much, much more than that? I think not.

The principle Lawford-Smith recommends is: 'You ought to do the least harm you can, compatible with a life worth living'. Think of this principle as giving you a sort of budget for avoiding harm, consisting of the resources you can spare whilst keeping enough for yourself to maintain a life worth living. The principle says you should use this budget to minimize the harm you do. To minimize harm, you should start by doing the things that most effectively reduce harm at a given cost. Offsetting your emissions is bound to be among those. It is such an effective means of reducing the harm you do that there should be no question of not taking it. If there are other things you can do that have such a favourable ratio of averted harm to cost, you should do those too.

Most of us will have offset all our emissions long before we have used up all our budget for minimizing harm. Offsetting is so cheap that even profligate emitters can offset their emissions for a few hundred dollars a year. What about people who find their budget is all gone before their emissions are fully offset? They are in a truly unfortunate position: in order to maintain a life worth living they must hang on to their last \$10, even at the cost of doing much, much more than £50 worth of harm to others. They must have very little money to spare in the first place. People in this position are among the poor, and the poor emit little greenhouse gas. For this reason, they are anyway under no duty of justice to offset their emissions.

Here I have to reiterate a point I made in *Climate Matters*,¹⁰ since Lawford-Smith seems to have missed it. The harm a person does by her emissions is an injustice only if it is not reciprocated. But the harm done by the meagre emissions of the world's poor is fully reciprocated: the poor emit little, which means they do less harm to others than others do to

them. So their emissions are not unjust.

I argued in *Climate Matters* that offsetting emissions, regarded as a way of doing good, is not one of the most effective ones. For a given cost, you could do much more good by supporting a charity that cures tuberculosis, for example.¹¹ It follows from what I have been saying that, for \$10 spent on an effective means of doing good, you can do very much more than \$50 worth of good. But Lawford-Smith and I are debating ways of avoiding doing harm, not ways of doing good. Offsetting is exceptionally effective at avoiding harm.

Lawford-Smith is right that avoiding harm can require us to make some comparisons. But our empirical situation is that, when we make these comparisons, it will certainly turn out that we should offset our emissions of greenhouse gas. The ratio between the cost of avoiding harm this way and the amount of harm that can be avoided is so great that there is no question we should avoid this harm. For this case, the simplistic principle of justice that we should not do harm is good enough.

All this is true of our present empirical situation only. At present very little is being done in the world to control emissions of greenhouse gas. That is why offsetting is so cheap. If the world begins to respond more adequately to climate change, opportunities for offsetting would become scarcer and the cost of offsetting will rise. Eventually other means of reducing harm may become competitive with it.

4 Bou-Habib

Paul Bou-Habib takes up the public morality of governments, which I applaud. Furthermore, I agree with many of his conclusions. For example, I agree entirely with his conclusion about the example he presents in section 4 at the end of his paper.

However, I disagree with him in one fundamental way. I divided morality roughly into two departments: justice and goodness. Bou-Habib thinks we need to recognize a third. But so far as I can see, his third department is actually just goodness; it is not a third department at all. Since I take public morality to be regulated largely by goodness, it is therefore no surprise that I agree with his particular conclusions.

Why does Bou-Habib think a third department is needed? Because of the non-identity effect. He argues in his section 3 that climate policy aimed at promoting goodness cannot give proper weight to the interests of those future people whose identity depends on the policy. If that were true, it would be disappointing to me, since a leading claim of *Climate Matters* is that climate policy must be aimed at promoting goodness precisely because this is the way to give proper weight to the interests of those future people whose identity depends on the policy. I argued that the morality of justice cannot do this, and that is why climate policy is regulated by the morality of goodness. But Bou-Habib thinks that the morality of goodness is subject to the same objection as I made to the morality of justice.

If were right, I would be surprised as well as disappointed. *Climate Matters* contains a wide review of public morality aimed at promoting goodness, covering several different topics. No doubt my particular conclusions are debatable in detail, but the general aim of promoting goodness seems to provide a good approach to the issues.

In taking this approach, I assume the goodness of a world is determined by the wellbeing of the people who live in it. That is to say, there is a function from individuals' wellbeing to the goodness of the world. I shall call this a 'value function'. I assume the value function is impartial between people: it does not depend on which particular person has which particular amount of wellbeing. Consequently, when the goodnesses of two worlds are compared, it makes no difference whether they contain the same or different people. Goodness is unaffected by the non-identity effect. My theory of goodness respects what Derek Parfit calls

‘the no-difference view’ – the view that the non-identity effect makes no difference.¹²

But Bou-Habib rejects the no-difference view for goodness. He divides goodness into two kinds, which he calls ‘person-affecting’ and ‘impersonal’ respectively, and the person-affecting kind is affected by identity. Actually, he never attaches the adjective ‘person-affecting’ to ‘goodness’, but always to ‘improvement’, and he never attaches ‘impersonal’ to ‘improvement’, but always to ‘goodness’. This confuses me, since the two properties are supposed to be parallel to each other. Moreover, for a reason I shall explain later, the term ‘impersonal’ is misleading. So I shall use ‘impartial’ instead of ‘impersonal’, and I shall apply both the adjectives ‘impartial’ and ‘person-affecting’ to ‘improvement’.

Given these innovations, Bou-Habib’s two kinds of improvement can be defined as follows. One of two alternative worlds is an impartial improvement on the other if it is better according to a correct impartial value function. One world is a person-affecting improvement on another if it is an improvement for those people who exist in both worlds.

For example, $(1, 2, \Omega)$ is an impartial improvement on $(1, \Omega, 1)$, but not a person-affecting improvement; it is not better for anyone. In the notation I use, a world is represented by a vector. Each place in the vector stands for a particular person. If the person exists in a world, her place in the vector that represents that world is occupied by a number that represents her wellbeing. If the person does not exist in a world, her place is occupied by an ‘ Ω ’ to mark her absence.

A person-affecting improvement and an impartial improvement may go in opposite directions. For example, compare the worlds $A = (4, 1, \Omega)$ and $B = (3, \Omega, 3)$. A is a person-affecting improvement on B since it is better for the first person. And for the sake of argument we may assume B is an impartial improvement on A since it has a greater total of wellbeing, more equally distributed. Bou-Habib’s view is that in a case of conflict like this, the person-affecting improvement dominates over the impartial improvement. It dominates in the sense that, in a choice between A and B , A should be chosen. Bou-Habib says that a person-affecting improvement always dominates over an impartial improvement, so long as the person-affecting improvement is not negligible.

That is why he rejects the morality of goodness described in *Climate Matters*. It is the morality of impartial goodness, and he believes impartial goodness counts for nothing. Any climate policy will be good or bad for some people. So whatever pair of options is being considered, one of them will be a person-affecting improvement on the other. That one should be chosen, whatever impartial goodness has to say. Impartial goodness is therefore irrelevant.

What supports this claim about dominance? The only evidence Bou-Habib offers for it is two examples in section 3. For the sake of assessing them, I shall make these examples a little more specific and render them into my notation. I shall take the second first.

The second example is the claim that, even people who think equality is valuable think it should not be achieved at the cost of ‘levelling down’. We should not achieve equality by making someone worse off without making anyone better off. For example, they think $A = (2, 2, 2, 3)$ should be chosen over $B = (2, 2, 2, 2)$. I take it for granted they are right about that.

A is a person-affecting improvement on B . The example would therefore support Bou-Habib’s claim of dominance if B were an impartial improvement on A . This would be so if B is better than A according to a correct impartial value function. But any reasonable value function respects the principle that, if one option is better for someone than another, and worse for no one, then it is better. (I call this ‘the principle of personal good’.¹³) So A is better than B according to any reasonable value function, including any reasonable impartial value function. B is therefore not an impartial improvement on A . The example does not support

Bou-Habib's case for dominance.

An impartial value function that ranks A above B may still be egalitarian. For example, it may rank $(2, 2, 2, 2)$ above $(1, 5, 2, 2)$. You can be impartially egalitarian without supporting levelling down.

Bou-Habib's first example is the claim that extra people should not be added to the world's population at the expense of a loss to an existing person. For example, $A = (10, 10, 10, 10, \Omega)$ should be chosen over $B = (9, 10, 10, 10, 10)$. For the moment let us suppose this claim is true.

A is a person-affecting improvement on B . So, if the claim is true, it would support Bou-Habib's case if B were an impartial improvement on A . But there are many impartial value functions that rank A above B . A very popular one is 'average utilitarianism', whose value function is the average of people's wellbeing. As it happens, average utilitarianism is not a correct value function.¹⁴ But there are also correct value functions that rank A above B . My book *Weighing Lives* argues that a correct value function is 'neutral-level utilitarianism' or 'critical-level utilitarianism', in which the value of a world is given by the total of the amounts by which each person's wellbeing exceeds a particular neutral level. If the neutral level is above 10, this function ranks A above B .

However, if the neutral level is below 9, neutral-level utilitarianism ranks B above A . Furthermore, whatever the neutral level, it will always be possible to create a parallel example in which neutral-level utilitarianism ranks B above A ; we have only to change the numbers to something about the neutral level. And I claim neutral-level utilitarianism is the correct impartial value function, so B may be an impartial improvement on A . A remains a person-affecting improvement on B . So does this support Bou-Habib's argument that a person-affecting improvement dominates?

It does not, because we cannot take it for granted that A should be chosen above B . According to Bou-Habib, 'we', the public, think it should. I agree with him about that. We are mostly gripped by a strong intuition that I call 'the intuition of neutrality'. It was nicely expressed by Jan Narveson in the words: 'We are in favour of making people happy, but neutral about making happy people.'¹⁵ We think that adding people to the population of the world is ethically neutral, so it should not be done at the expense of existing people. However, philosophers have by now worked hard on population ethics for more than three decades. One conclusion that has emerged rather clearly is that we cannot assume all our natural intuitions about this subject are correct. Our natural intuitions turn out to contradict each other.¹⁶ In particular, I think we have to accept that the intuition of neutrality is generally false.¹⁷

At any rate, Bou-Habib should not rely on such an insecure intuition in order to reject the morality of goodness, along with the analysis of goodness in my book. Since it rests on just these two examples, and neither is convincing, his case against the morality of goodness is weak.

Two final notes. First, the idea of a person-affecting improvement is anyway unsatisfactory. Take this sequence of worlds: $(2, \Omega)$, $(3, 1)$, $(\Omega, 2)$, $(1, 3)$. Each is a person-affecting improvement on the one before, and the first is a person-affecting improvement on the last. So the relation 'is a person-affecting improvement on' is cyclical. But an improvement is an increase in goodness, so it is an analytical truth that a genuine relation of improvement cannot be cyclical. Furthermore, this makes it clear there can be no such thing as person-affecting goodness. Perhaps this is why Bou-Habib avoids referring to it.

Second, the term 'impersonal goodness', which Bou-Habib adopts along with many other authors, is very misleading. He may find it rhetorically effective in diminishing the

importance of what it denotes: how could anyone support impersonal goodness against the good of persons? But actually what he calls ‘impersonal goodness’ is nothing other than the good of persons. A more accurate term for it would be ‘personal goodness’. A plausible objection to the treatment of goodness in *Climate Matters* would be that it is too much focussed on personal goodness and not enough on animal goodness or natural goodness. Perhaps I should have taken more notice of genuinely impersonal goodness.

References

- Arrhenius, Gustaf, (forthcoming) *Population Ethics: the Challenge of Future Generations*, Oxford University Press.
- Broome, John, (1991) *Weighing Goods: Equality, Uncertainty and Time*, Blackwell.
- Broome, John, (2004) *Weighing Lives*, Oxford University Press.
- Broome, John, (2012) *Climate Matters*, Norton.
- Cullity, Garrett, (2015) ‘Acts, omissions, emissions’, in *Climate Change and Justice*, edited by Jeremy Moss, Cambridge University Press, 148–64.
- Narveson, Jan, (1973) ‘Moral problems of population’, *The Monist*, 57 (1973), pp. 62–86.
- Intergovernmental Panel on Climate Change Working Group 3, (2014) *Climate Change 2014: Mitigation of Climate Change*, Cambridge University Press.
- Parfit, Derek, (1984) *Reasons and Persons*, Oxford University Press.
- Sinnott-Armstrong, Walter, (2005) ‘It’s not my fault: global warming and individual moral obligations’, in *Perspectives on Climate Change: Science, Economics, Politics, Ethics. Advances in the Economics of Environmental Research*, Volume 5, edited by Walter Sinnott-Armstrong and R. B. Howarth, Elsevier, pp. 293–315.
- Scanlon, T. M., (1998) *What We Owe to Each Other*, Harvard University Press, 1998.

Notes

My research for this paper was supported by the Australian Research Council (ARC) Discovery grant number DP140102468.

1. *CM* pp. 73–4.
2. For example, Cullity (2015) and Sinnott-Armstrong (2005).
3. *CM* p. 66.
4. *CM*, p. 79.
5. For example, Scanlon (1998)
6. *Reasons and Persons*, pp. 75–82.
7. US Government Interagency Working Group on Social Cost of Carbon, *Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis*, May 2013 revised July 2015. The report gives figures in 2007 US \$s, which I have converted to 2016 US \$s.
8. *CM*, p.74.
9. IPCC, *Mitigation of Climate Change*, section 3.6.2, p. 229.
10. *CM*, p. 58.
11. *CM*, pp. 65–6.
12. *Reasons and Persons*, p. 367.
13. *Weighing Goods*, chapter 8.
14. As is explained in my *Weighing Lives*, section 13.3.
15. ‘Moral problems of population’.
16. See the ‘impossibility theorems’ in Arrhenius, *Population Ethics*.
17. *Weighing Goods*, pp. 205–8.