

Climate change and optimum population

Hilary Greaves

July 5, 2017

1 Introduction

This paper is a contribution to the debate over optimum human population. More precisely, the topic of debate is the optimal medium-term instantaneous population size: the optimal number of human beings for there to be living on Earth at any one time, over the next few decades or so.

The considerations relevant to this question are many and various, and I will not attempt anything like a comprehensive survey of them here.¹ But one highly relevant set of considerations relates to anthropogenic climate change. The paper attempts to clarify whether, and if so precisely why, climate change generates reasons for favouring a reduction in instantaneous population size, relative to natural choices of status quo. I will argue that the usual arguments for this conclusion are too simplistic: in particular, they ignore the fact that (in a precisely specifiable sense) climate change depends on *cumulative* emissions across all time, rather than on emissions *rates* per se. The same conclusion can plausibly be supported once that fact is accounted for, but the relevant arguments are more complex (and less conclusive).

The paper proceeds as follows. Section 2 introduces the usual simple argument for the claim that, because of climate change, a lower instantaneous population would be better. This argument is based on the very natural thought that since emissions are made by people, having fewer people would naturally lead to lower emissions. The argument, however, essentially gives centre stage to the notion of emissions *rates*. Section 3 notes the sense in which climate change depends primarily on cumulative emissions, rather than on emissions rates. Section 4 introduces a way of conceptualising what we would be doing if we reduced instantaneous population sizes, in terms of spreading out a given number of lives over a longer stretch of time (what I will call ‘population spreading’). By

¹I attempt a partial survey in (Greaves, 2018).

thinking in this way, we see that it is far from obvious that population reduction would reduce cumulative emissions, and hence (given the points of section 3) far from obvious that it would in the end generate any improvement. Using the notion of population spreading, I set out a simple model of what the effects of population reduction might be such that (I claim) *if* this simple model were accurate, reducing instantaneous population sizes would in fact not help with climate change at all.

In section 5, I survey two possible reasons one might have for doubting this conditional claim: for thinking, that is, that even if the simple model were accurate, still population reduction would help with climate change. In each case, I argue that the reasons in question are confused (either because they rely essentially on a false theory of value, or because they do not in the end have the claimed implication). Section 6 turns to the question of whether the assumptions of the simple model are correct. I argue there that two of the simple model's assumptions are likely to be importantly inaccurate, and that because of this, population reduction probably would, in the end, help with climate change. The key issues concern (i) ways in which population reduction might reduce even *cumulative* emissions, and (ii) ways in which merely slowing down the course of climate change, even if it did not correspond to reducing the extent of the climate change that eventually occurs, could be beneficial (because of the nature of various possible processes of adaptation). The aim of the paper is to focus attention and debate onto these key issues, and away from the too-simplistic idea that “fewer people means less greenhouse gas emission”.

2 The emissions-rate argument for population reduction

Climate change, as is well known, is driven by emissions of various greenhouse gases (GHGs). And GHG emissions rates are on the increase. Prior to the industrial revolution, emissions were (relatively speaking) extremely low; by 1970 the total global emissions rate had reached the equivalent² of 27 billion tons of CO₂ annually; by 2010 the rate had reached 49 billion tons per year,

²There are many different greenhouse gases — carbon dioxide is the biggest contributor to climate change, but by no means the only contributor — and these gases differ from one another over (crucially) their time-profiles of contribution to radiative forcing per unit mass of gas. For example, carbon dioxide is relatively long-lived in the atmosphere, and consequently contributes very significantly to radiative forcing over a timescale of hundreds of years; methane contributes more forcing per unit mass initially, but is much shorter-lived in the atmosphere. As a result, there is no single canonical way of representing the relevant multidimensional facts on a single real-valued scale (that is, no single canonical ‘emissions metric’). The figures quoted in the main text follow the IPCC Fifth Assessment Report in using the ‘GWP₁₀₀’ emissions metric, which quantifies emissions based on their cumulative contributions to radiative forcing over the first 100-year period (IPCC, 2014b, box TS.5).

and this emissions growth is expected to persist for some time (IPCC, 2014b, pp.45-49). As a result of these emissions, climate models predict, the world is likely to experience a surface warming of between about 1.5 and 4°C by 2100, and possibly further significant increases beyond that date. The early stages of this warming have already been observed (IPCC, 2013, esp. pp.37, 89).

What causes this rise in emissions rates? Commentators vary somewhat in their tendency to point the finger of blame at increasing average GDP per capita, or instead at increasing population sizes. Both are of course contributors, at least in the purely arithmetical sense that the total emissions rate is given by (average GDP per capita) \times (population size) \times (emissions intensity of GDP). (This equation is a condensed form of the Kaya identity; see e.g. (IPCC, 2014b, box 5.1).) And, at the global level³, both of the first two factors in this product *are* increasing.⁴ Over the period 1970-2010, for instance, the global average GDP per capita rose by 103%, and the global population size rose by 87% (IPCC, 2014b, p.365); both are expected to continue rising (at least) well into the present century ((IPCC, 2014b, figure TS.7),(Nations, 2017, figure 2)). For present purposes, though, we need not single out one of these factors as ‘the culprit’ to the exclusion of the other: holding fixed the remaining factors in the Kaya identity, a given proportional increase in either equally increases the total emissions rate by the same proportion. Insofar as we have a goal of reducing the total emissions rate, then, reducing population size looks like at least one sensible thing to try.

This general line of thought is outlined by (among others) Jonathan Porritt, in his ‘call to arms’ for population reduction^{5,6}:

³There are of course significant regional variations in each of the three factors. But consideration of the global averages will suffice for the purposes of this paper.

⁴The (global average) emissions intensity of GDP, meanwhile, is decreasing. This offsets the increases in GDP per capita and in population size to some degree, but not enough to prevent the total global emissions rate from increasing IPCC (2014b, p.365).

⁵I focus on Porritt’s version of the argument for concreteness and because Porritt sets the argument out more clearly than most, but the basic line of thought is often heard. In particular, my guess is that this is what many authors have in mind (whether or not the authors are willing explicitly to draw conclusions about the desirability of population control) when they claim that that population growth, via its contribution to rising emissions rates, is one of the key drivers of climate change (see e.g. (IPCC, 2014b),(Royal Society, 2012, p.68),(Population Matters, 2016)).

⁶Here and throughout, the relevant sense of ‘population reduction’ is counterfactual rather than temporal. For instance: according to the United Nations’ ‘medium’ population projection, population will grow from 7.3 billion in 2015 to 9.7 billion in 2050. According to the ‘low’ projection, meanwhile, population in 2050 will rise to only 8.7 billion (United Nations, 2015). The ‘low’ scenario counts as population reduction relative to the medium scenario, notwithstanding the fact that even according to the low scenario, population in 2050 is higher than in 2015.

Porritt’s reference to a ‘downward population trajectory’ (quoted below) seems to be temporal in the first instance, but since the implied contrast is presumably a state of affairs in which the population trajectory is not a downward one, the counterfactual interpretation is valid here too.

At one level, this is all about basic mathematics. We roughly know the total volume of greenhouse gases we can put into the atmosphere over the next few decades if we are to stay the right side of the two-degree-centigrade increase (by the end of the century) which scientists tell us we absolutely mustn't go above. That total volume has to be divided up between the total number of people doing the emitting... I believe it's part of our duty to the next generation... to advance the compassionate, progressive case for a full-on global campaign to put the world on a downward population trajectory just as fast as we can.(Porritt, 2010)

In fact, Porritt's argument is open to two quite different interpretations. Both start from the crucial claim that, for some fixed R , *in order to avert dangerous climate change we need to stay below total emission rate R* . Thereafter, the two interpretations diverge. The first emphasises that life over the next few decades can be *more pleasant* given a lower population, consistently with averting dangerous climate change: the total emission rate R 'has to be divided up between the total number of people doing the emitting', so that we each get a larger share (and thus arguably a better life) if there are fewer others around for us to share with. The second emphasises instead that it is *not a foregone conclusion we will succeed* in averting dangerous climate change, and that this becomes more likely under lower-population scenarios (due to the difficulties involved in significantly reducing *per capita* emissions).

For either version of Porritt's argument to work, however, it is essential that the key to averting dangerous climate change is reducing the emissions *rate* (or, more or less equivalently, reducing total emissions within some *fixed* number of decades: in particular, it must be the same number of decades that is relevant across scenarios that differ from one another over the path of instantaneous population size). Porritt and his fellow-travellers seem to be thinking that the crucial period for emissions is, say, between now and 2050, regardless of how many people live during that time. But in fact the *timing* of emissions has only a rather limited significance. To see this, we need a brief excursion into the relevant climate science.⁷

⁷One might try to be more charitable: perhaps Porritt is relying on fact that the 'carbon budget' that we must stay within, for the purpose of averting dangerous climate change, will be used up within a few decades whether or not we achieve population reduction. Given that fact, it is indeed the case that 'we roughly know the total volume of greenhouse gases we can put into the atmosphere over the next few decades' if we are to avert dangerous climate change. However — as I will argue in this paper — unless this translates somehow into a fixed cap on the total emissions rate, it takes a lot more than 'basic mathematics' to establish that 'putting the world on a downward population trajectory' would help.

Similarly, one might be interested only in emissions within the next few decades if one thought that due to exogenous technical progress, emissions will be effectively zero after (say) 2070 in any case. But technical progress is due to work carried out by people; it is therefore implausible that the rate of technical progress is exogenous in the relevant sense (i.e., independent of instantaneous population sizes). I investigate the effect of population size on technical progress rate in (slightly) more depth in section 6.1.

3 Cumulative emissions

An *emissions pathway* is an assignment of total global emissions rates to times. Climate science seeks, among other things, to predict which emissions pathways would lead to which *temperature pathways*: assignments of global mean surface temperatures to times. Given realistic climate science, together with the fact that total cumulative emissions are limited by the amount of fossil fuel in the ground, one expects that for any possible emissions pathway, the temperature pathway will consist of ascent to a peak, followed by (eventually) a decline.

The key finding from climate science ((Matthews & Caldeira, 2008; Allen *et al.*, 2009; Eby *et al.*, 2009; Matthews *et al.*, 2009; Zickfeld *et al.*, 2009)) is that⁸

(CE-PH) Peak warming is, to a very good approximation, determined by cumulative emissions alone, i.e. independently of any other details of the emissions pathway.

The observation (CE-PH) is highly significant for policy purposes. It means that merely *slowing down* the rate at which we (collectively) burn fossil fuels will not be enough to avert dangerous climate change; what we crucially need is eventually to get to the point of *near-zero* carbon emissions (so that cumulative emissions are finite), and to do so while significant fossil fuel reserves remain in the ground (so that cumulative emissions are sufficiently small to limit climate change to a reasonable extent).⁹ Thus it suggests that policies should ultimately be favoured on the basis of the extent (if any) to which they increase the probability of this desirable eventual outcome; effects on emissions *rates* are of interest only insofar as they are plausibly related to that ultimate goal.

As a special case of this, (CE-PH) is highly significant to the optimum population debate. For, as I will elaborate in the next section, it suggests that

⁸The focus on *peak* warming is particularly prominent in (Allen *et al.*, 2009). What the literature establishes more generally is that a broad range of climate responses to emissions (including not only the peaks but also the whole pathways for atmospheric CO₂ concentration, global mean surface temperature, ocean acidity, and other relevant variables) depend on cumulative emissions alone, and are in this sense path-independent, in the long term. (Of course there is path-dependence in the short term: for example, holding fixed cumulative emissions, the number of degrees of warming experienced in 2050 obviously depends on how much of that fixed total is emitted before vs. after 2050.) This generality is relevant for the present paper, since it implies that some ways of caring about variables other than peak warming will fail to generate any reason to care about instantaneous population size. Nevertheless, in this paper I focus on the claim (CE-PH), for simplicity of exposition.

This path-independence is perhaps unsurprising, given the very long lifetime of CO₂ in the atmosphere ((Archer, 2005)), and the fact that the literature is only considering variations to the emissions path on a much shorter timescale (tens or hundreds of years).

⁹Negative net carbon emissions could in principle be achieved by widespread use of ‘carbon capture and storage’ technology, driven by non-carbon energy sources. This means, *inter alia*, that an eventual return to low cumulative net emissions is in principle consistent with burning all the fossil fuel. I leave this complication aside for simplicity of exposition.

(say) halving instantaneous population might well not help at all with climate change *if* it merely resulted in us taking twice as long to emit the same cumulative emissions. *Pace* the emissions-rate argument, (CE-PH) implies that insofar as our concern is to limit peak warming, our concern should not be with the total amount of carbon we emit *over the next few decades*, but rather the total amount of carbon we emit *ever*. (Relatedly, it is unclear why we should be fixated specifically on avoiding a temperature increase of more than 2 °C *by the end of the century*, rather than seeking to limit the height of the warming peak to (say) 2°C *whenever it occurs*.) To put it another way: it is unclear whether there is any value to reducing the global emissions *rate*, if this is done via population reduction.

4 Population reduction as population spreading

The purpose of the present section is to introduce a concept that will prove helpful for clarifying the question of whether and how population reduction might help with climate change, given the pre-eminent significance of cumulative emissions. This is the concept of *population spreading* (more precisely — a refinement I will introduce in section 4.3 — *mere emissions-era population spreading*).

4.1 Population spreading

Let A be any state of affairs (in the timeless sense, i.e. a complete history of the universe, from big bang to heat death). The specification of A includes information about how many individuals are alive at each time and their per-capita emissions levels at each time, and thus determines (*inter alia*) a particular emissions pathway and thence a particular temperature pathway.

Let us now define a second state of affairs B, which differs from A as follows. There is a bijection¹⁰ π from the set of persons who ever exist in A to the set of persons who ever exist in B, such that:

- π preserves life lengths. That is, for every person x who exists in A, the life that x lives in A and the life that $\pi(x)$ lives in B are equally long.
- π respects the birth order. That is, for all persons x, y who ever exist in A, if x is born before y in A, then $\pi(x)$ is born before $\pi(y)$ in B.

¹⁰We of course expect that the *identities* of persons in B are different from those in A; but we need not dwell on this fact, because we assume that bare identities of persons are morally irrelevant for present purposes.

- There is some fixed date D (e.g., 2017) such that (i) A and B are qualitatively identical before D ¹¹, and (ii) for all individuals x, y who are born after D in A , the interval between the births of $\pi(x)$ and $\pi(y)$ in B is longer than the interval between the births of x and y in A .

When these conditions obtain, say that B is related to A by *population spreading*. Heuristically, the idea is that what one does, in ‘moving’ from A to B , is to spread people out more thinly across the decades. Thus the *instantaneous* population at any given time (after D) is smaller in B than in A , but the *timeless* population (the set of people who ever exist) is the same. There is one (instantaneous) sense in which a transition from A to B involves reducing population size, and another (timeless) sense in which it does not.

Now suppose further that

- π preserves individual lifetime carbon emissions. That is, for all persons x who exist in A , the lifetime carbon emissions of x in A are equal to the lifetime carbon emissions of $\pi(x)$ in B .

When this further condition also obtains, say that B is related to A by *mere population spreading*. Note that whenever this relation obtains, A and B have the same cumulative emissions as one another. Although B in one sense has a lower population than A , *this* way of reducing population size would not reduce cumulative emissions.¹²

4.2 A simple model

Consider the following three conditions:

(MPS) Reducing instantaneous population size would implement mere population spreading.

(CE-PH) Peak warming is determined by cumulative emissions alone.

(PH-CI) Climate-change impacts are a function of peak warming alone.

¹¹It follows (given our assumption that π respects the birth order) that π is the identity map before D .

¹²Note also that the conditions specified here are consistent with the time-profile of instantaneous population size in A being e.g. as in the UN medium projection, and in B being as in the UN low projection.

I do not *assert* these conditions; section 6 will examine how plausible each of them is (and, indeed, will argue against the first and third conditions). But, as a step towards clarifying the debate, it is useful to consider what follows *on the assumption that* these conditions all hold.

What follows from the conditions (MPS), (CE-PH) and (PH-CI) jointly — since mere population spreading does not affect cumulative emissions — is that reducing instantaneous population size would not reduce climate-change impacts¹³, and in this sense would not ‘help with climate change’.¹⁴ Although B has a lower instantaneous population and consequently a lower emissions *rate* than A, *this* way of achieving a reduction in emissions rate is not one that considerations of climate change would give us any reason to welcome.

4.3 Emissions-era population spreading

It is worth making one refinement to the simple model of section 4.2. The issue is that as I have defined ‘mere population spreading’, that concept applies only when the ‘spreading’ in question does not affect the number of people who ever live (it does not, that is, affect the ‘timeless population size’). However, (1) it is very implausible that a reduction in instantaneous population size would not affect the timeless population size¹⁵, and (2) this feature is anyway inessential to the simple model. In this section, I will therefore define a weaker notion of population spreading, which does not have this unwanted condition built into the definition.

Given the finitude of fossil fuel reserves, there will eventually come a date beyond which no significant further net carbon emissions will be made. Let t_{end} be such

¹³It is not immediately clear precisely what sort of thing counts as an ‘impact’; I return to this in section 5.

¹⁴This is of course not to say that B is not better overall than A: climate change is not the only influence on overall value. In fact, I have not specified the relationship between A and B precisely enough to determine which, if either, is better: it is consistent with my specification, for instance, that there is a major world war in the 22nd century A but not in B, or vice versa. I am assuming only that a contribution from ‘climate change impacts’ can be meaningfully separated from any unrelated other ways in which A and B may differ. More importantly, I am also assuming that ‘climate change impacts’ can be meaningfully separated from other value-relevant ways *caused by population spreading* in which A and B differ. The latter question is of course relevant to the question of whether population spreading would make things better or worse overall; I return to it briefly in the concluding section of this paper.

¹⁵It is unclear whether a short-to-medium-term reduction in instantaneous population size would increase, or instead decrease, the size of the timeless population; there are plausible considerations pointing in each direction. I return to this matter briefly in section 7. But it would take an inexplicable coincidence for the effect (even: the effect on *expected* timeless population size) to be zero.

a date. Let us refer to the period from the present¹⁶ to t_{end} as the ‘emissions era’.

Again, let A be any ‘status quo’ state of affairs. Now let B be related to A by spreading people out more thinly across time *during the emissions era*, but let us assume nothing about how B is related to A, in particular as regards population sizes, at later times. That is, suppose that there is some bijection π from a subset of A’s (timeless) population to a subset of B’s (timeless) population, such that

- All persons who exist during the emissions era in either A or B are in (respectively) the domain or range of π .
- π respects the birth order.
- π increases birth spacing.

When these conditions obtain, say that B is related to A by *emissions-era population spreading*. If (further) π preserves lifetime carbon emissions, say that B is related to A by *mere emissions-era population spreading*.

The point is then that the simple model of section 4.2 works just as well if the condition **(MPS)** is replaced by the weaker condition **(MPS-EE)**:

(MPS-EE): Reducing instantaneous population size would implement mere emissions-era population spreading.

Since it follows from **(MPS-EE)** that reducing instantaneous population size would (to a very close approximation) leave cumulative emissions unaffected, the claim that reducing instantaneous population size would not reduce climate impacts goes through (to a very close approximation) just as before.

4.4 Physical time and population time

The following conceptual apparatus will be helpful to the discussion. Let us distinguish between *physical* time on the one hand, and what we might call *population time* on the other. Physical time is the time we are all familiar with, and that is measured by standard clocks: one physical year passes each time the Earth completes one revolution around the Sun, one physical second passes

¹⁶A more principled definition would identify, say, the start of the industrial revolution, rather than the present, as the start of the ‘emissions era’. However, for present decision purposes there is no point in considering counterfactual variations to the past, thus defining the ‘emissions era’ as beginning at the present moment will serve our present purposes.

while a light ray travels 3×10^8 metres in a vacuum, and so on. Population time – I hereby stipulate – is a modification of physical time, where the modification is based on instantaneous population size: let the number of seconds of population time that pass between any two given instants be equal (by definition) to the number of physical seconds elapsed *multiplied by* the instantaneous population size during the interval in question.¹⁷ Thus, for example, if the instantaneous population is smaller in state of affairs B than in state of affairs A, then the rate at which the population clock ticks (relative to the physical clock) is lower in state of affairs A than in state of affairs B.

This notion of population time will allow us to frame the issues in more illuminating ways. For example, Porritt is quick to note the obvious fact that reducing instantaneous population size (while holding fixed per-capita emissions) reduces the total global emissions rate *relative to physical time*. But for all that, it could easily have no effect on the total global emissions rate *relative to population time*. One question, therefore, is which of these rates (if either) is positively correlated with expected cumulative emissions, or is via any other route positively correlated with expected climate impacts. And this is an open question: nothing that we have surveyed so far gives us any reason to think that the physical-time rate is more relevant than the population-time rate in either of these senses. Certainly, nothing in Porritt argument itself even begins to address the question.

To summarise and preview: using the notion of (mere emissions-era) population spreading (sections 4.1 and 4.3), I have set out a simple model of what the consequences of reducing population size could conceivably be (section 4.2), and made the conditional claim that *if* the assumptions of that simple model were correct, then reducing population size would not help with climate change. The next section examines and defends the conditional claim, while section 6 investigates the plausibility of the model’s assumptions.

5 Intrinsic reasons for favouring a physical-time slowdown?

I noted in section 4 that as a matter of logic, it follows from the assumptions of the simple model that reducing instantaneous population size would not ‘reduce climate impacts’. But the term ‘climate impacts’ has so far been merely a placeholder.

¹⁷More precisely: $dt_{pop} := N(t) \cdot dt_{phys}$, where t_{phys} is physical time, t_{pop} is population time and, as before, $N(t)$ is instantaneous population size at t . (‘The’ instantaneous population size during an *interval* of time is of course not in general well-defined, so the definition-sketch given in the main text does not strictly speaking make sense.)

We might most naturally take ‘impacts’ to be medium-scale physical changes that affect the quality of individual lives: changes in, for example, sea level, frequencies of droughts and floods, and the geographical distribution of various diseases.¹⁸ If so, then ‘impacts’ are, roughly, things that affect the average well-being of persons who live at the time and place at which the impact in question occurs. And in that case, it does not immediately follow that reducing population size would not *make things better* for reasons related to climate change. We should distinguish, conceptually, between ‘would not reduce climate impacts’ and ‘would not generate any climate-related improvement’.

One way in which these two notions can come apart involves axiologies that, intuitively, assign intrinsic importance to matters of physical time: axiologies according to which it matters not only how many lives are lived and how good those lives are, but also *when* these lives (and components of these lives) are lived, relative to the physical clock. In this section, I will briefly survey two such axiologies, and the way in which each could be used to argue that even if climate impacts (in the above sense) were unaffected, still the results of population spreading would count as making things better.

Both arguments are unsound. I survey these arguments only to set them aside, and to move on to more promising arguments in section 6.

5.1 First bad argument: via Instantaneous Averagism

5.1.1 Instantaneous averagism

The first argument I will consider proceeds by assuming a particular axiology, which I will call ‘instantaneous averagism’. This takes a bit of setting up.

Ultimately, we normally want to work with axiologies that take as input the data on how well each person’s life goes *as a whole* — the person’s *lifetime* well-being level — and that output an index of the overall goodness of the

¹⁸This appears to be the way the IPCC uses the term ‘impacts’. The IPCC’s official account is as follows:

In this report, the term *impacts* is used primarily to refer to the effects on natural and human systems of extreme weather and climate events and climate change. Impacts generally refer to effects on lives, livelihoods, health, ecosystems, economies, societies, cultures, services, and infrastructure due to the interaction of climate changes or hazardous climate events occurring within a specified time period and the vulnerability of an exposed society or system. IPCC (2014a, p. 5)

The validity of the argument from (1) the simple model’s assumptions to (2) the conclusion that reducing population size would not reduce climate impacts is of course independent of the interpretation of ‘impacts’, provided that the interpretation is the same in the premises and the conclusion.

state of affairs.¹⁹ In the context of real-world applications, however, it is often convenient to work in the first instance with data about *instantaneous* well-being levels, corresponding to the matter of how well people’s lives are going *at a particular instant*. Correspondingly, we might seek (in the first instance) an account of how good a given state of affairs is *at a particular moment in time*. On at least one natural approach²⁰, the answer to that question is determined by an *average* of how well individuals’ lives are going at the time in question. In symbols, the quantity in question is

$$\frac{1}{|P(x, t)|} \sum_{i \in P(x, t)} w_i^{inst}(x, t)$$

where, for each state of affairs x and time t , $P(x, t)$ is the set of persons who exist at time t in state of affairs x , and $w_i^{inst}(x, t)$ is the instantaneous well-being level that person i enjoys at time t in state of affairs x .

Even if we are working with such ‘instantaneous axiologies’, for decision purposes we ultimately need to compare timeless states of affairs.²¹ Thus, an instantaneous axiology is decision-relevant only insofar as it is related to a timeless axiology. One natural proposal for such a relation takes the goodness of a timeless state of affairs to be represented by the sum of its instantaneous goodness levels, where the summation is carried out across all times²²:

$$V_{IA}(x) = \sum_t \frac{1}{|P(x, t)|} \sum_{i \in P(x, t)} w_i^{inst}(x, t)$$

Thus we arrive at the axiology of interest: goodness is equal to the sum-over-times of average instantaneous well-being. Call this axiology *instantaneous averagingism*.²³

¹⁹Or, at least: an index of how good the state of affairs is *vis-a-vis considerations of welfare*. *Welfarism* is the view that goodness is entirely determined by welfare, but one need not believe welfarism to think that an index of welfarist goodness is at least one important input to assessments of overall goodness. Here, I consider only welfarist goodness, for simplicity (thus I don’t consider, for example, any evaluative significance that might attach to intact ecosystems, over and above their contributions to welfare).

²⁰This is of course not the only possible approach. For instance, one might well consider the *total* instantaneous well-being at each point in time, rather than the average. I return to this possibility below (section 5.2).

²¹To base decisions at t only a comparison of alternative possible states of affairs *at t* would be to ignore the effects of our actions on goings-on at later times; and once we seek to take the latter into account, we will face questions of trade-offs between benefits and costs felt at different times. A method for resolving these tradeoffs generates comparisons of timeless states of affairs.

²²More precisely—since time is continuous—by the time-integral of instantaneous goodness: $V_{IA}(x) = \int dt \cdot \frac{1}{|P(x, t)|} \sum_{i \in P(x, t)} w_i^{inst}(x, t)$.

²³Hurka (1982a) discusses this axiology under the label ‘A2’.

If we assume instantaneous averaging, we might reason as follows. Even if merely slowing down climate change relative to (physical) time fails to reduce ‘climate impacts’ *in the sense that* for every damaging climate-change related event (every storm, or flood, or drought, etc.) in A, a precisely similar event occurs at some time in B, the fact that these damaging events generally occur at *later* times in B amounts to an improvement. That (we might argue) is because what we do, in spreading out the process of climate change over a longer stretch of time, is increase the length of time during which humanity is able to enjoy the relatively benign climate conditions that precede the warming peak, correspondingly reducing the length of time during which humanity has to suffer the relatively blighted post-peak climate conditions.

As it stands, this argument is very rough. Its basic idea could be made precise in various ways. However, we need not concern ourselves with the details of that task; the important thing to notice is that *the instantaneous-averagist axiology is completely implausible*, so that any argument that relies essentially on instantaneous averaging is anyway unsound. The next section argues for this implausibility claim.

5.1.2 Against instantaneous averaging

In fact, averagist population axiologies are implausible in general. For example, the timeless averagist axiology, according to which the overall goodness of a (timeless) state of affairs is given by the average of the lifetime well-being levels of all persons who ever exist in that state of affairs, itself has deep problems²⁴, which are shared by instantaneous averaging’s way of valuing instantaneous states of affairs. However, there are also deeper problems with instantaneous averaging.

The deeper problems arise from the fact that according to this axiology, welfare goodness fails to supervene on the assignment of *lifetime* well-being levels to persons. Such a supervenience assumption is common ground to every population axiology that has been seriously debated by moral philosophers (including

²⁴This is one of the few things on which just about all population axiologists agree. For example, any form of averaging implies a corresponding ‘Sadistic Conclusion’, according to which it can be better to add a smaller number of persons with negative well-being than to ‘add’ a larger number of persons with positive well-being (to a common ‘pre-existing’ population whose members are to be unaffected by any such addition). Perhaps still worse, according to averaging it can be better to add people with negative well-being than not to add them, provided only that the pre-existing average is still more negative. This is not to say that any counterintuitive implication immediately grounds a fatal objection to a population axiology; that cannot be so, since it is well-known that all population axiologies have counterintuitive implications (Parfit, 1984; Ng, 1989; Carlson, 1998; Arrhenius, 2000, n.d.). But the implications of averaging are more intuitively unacceptable than most. For more on these and other criticisms of averaging, see (Hurka, 1982a,b; Arrhenius, n.d.).

timeless averagism). The assumption means, for instance, that if Joe suffers extra misery in his youth but also extra satisfaction in his middle age, and that these two changes balance each other out in the sense that Joe's lifetime well-being is unaffected, and in addition no-one else's existence or lifetime well-being is affected by the change, then overall goodness is unaffected. In contrast, according to the instantaneous-averagist axiology V_{IA} , such a change could easily affect overall value. If the instantaneous population is smaller (resp. larger) at the time of Joe's youth than in his middle age, then V_{IA} will attach more weight (resp. less weight) to increments to Joe's well-being in youth than in middle age, so that the change in question is likely to decrease (resp. increase) overall value according to V_{IA} .²⁵

The supervenience principle discussed in the preceding paragraph is the Pareto Indifference principle: if states of affairs A and B are equally good for each person, then they are also equally good *simpliciter*.²⁶ Closely related to the fact that V_{IA} violates the Pareto Indifference principle is fact that it also violates other (fixed-population) Pareto principles. For example, it violates the Strong Pareto principle, according to which, if B is better than A for some individual and worse for no-one, then B is better than A.

The most flagrant way in which V_{IA} violates the Strong Pareto principle concerns premature death. According to V_{IA} , it would constitute an improvement if anyone whose instantaneous well-being is consistently below average died sooner rather than later, despite the fact that this would (in most cases) reduce the lifetime well-being of the victims, while (we may stipulate) it would not increase the well-being of those who live on. If, for example, at every time there are some persons at each of the instantaneous well-being levels 100, 101, 102, . . . , 110, and if in addition a given person's well-being is constant throughout her life, V_{IA} implies that it would be better if everyone with well-being levels below 110 died as soon as possible. Particularly when (as we may further stipulate) a well-being level of 100 corresponds to a life going extremely well, so that premature death (presumably) constitutes a large reduction in lifetime well-being for the victim, this is completely implausible.

The Pareto principles for fixed population sizes are among the least controversial principles of moral philosophy. Even someone who is willing to accept the

²⁵Whether this 'likely' condition holds depends on what the relationship is between momentary and lifetime well-being. The condition stated in the main text follows if lifetime well-being is simply the time-integral of momentary well-being within the life in question. On other accounts of that relationship, it could happen that some change (affecting only two 'times') leaves Joe's *lifetime* well-being unaffected, but increases his momentary well-being at one time far more than it decreases it at the other affected time, in which case the condition stated in the main text need not follow.

²⁶I interpret the Pareto principles in terms of betterness, rather than in terms of preferences. The preference interpretation is more common (in particular, in the economics literature), but is irrelevant to the present evaluative enterprise, except insofar as the resulting principle collapses into the betterness-based principle (via assuming a preference-satisfaction theory of well-being). Broome (1991, p.165) coins the term 'Principle of Personal Good' for the 'betterness version' of the Strong Pareto Principle.

drawbacks of the timeless averagist axiology should not be willing to accept these violations of Pareto principles, and thus should not accept instantaneous averagism.

This is important, since it is fairly common, in the multidisciplinary environment of the optimum population discussion, for authors implicitly to rely on something like instantaneous-averagist assumptions. We can disregard any argument that relies essentially on an instantaneous-averagist axiology. In particular, we can disregard the argument sketched in section 5.1.1.

It is worth also noting the following more general point. The standard Pareto principles implicitly assume a fixed population. We can generalise slightly by requiring merely that the states of affairs being compared contain populations of the same size, and formulating ‘anonymised’ Pareto principles that disregard any information about personal identities. (For example, the anonymised Pareto Indifference principle says that if there is a bijection π from the population in x to the population in y such that for all $i \in P(x)$, i ’s well-being in x and $\pi(i)$ ’s well-being in y are equal, then x and y are equally good.) By definition, then, *any* axiology that respects the anonymised Pareto Indifference principle will judge a transition from A to B to be neither an improvement nor a disimprovement, if there is a bijection π between the set of people (ever) existing in A and the set of people (ever) existing in B, *and if in addition this bijection preserves lifetime well-being levels*. Since such bijections can easily fail to preserve the times at which people live, however, any axiology that assigns intrinsic importance to time will judge B better than A in some such cases, and will therefore violate the anonymised Pareto Indifference principle. It will for that reason be an implausible axiology.²⁷ Despite this, I turn next to investigating a second axiology that assigns intrinsic value to physical time, since it too is quite popular.

5.2 Second bad argument: Pure discounting

The second axiology (or family of axiologies) I will consider holds that a given change in well-being (say, due to climate-change impacts) simply *matters less*, the further in the future it occurs. Call this the assumption of pure discounting.

Given the assumption of pure discounting, one might try to reason as follows: even if population spreading does not affect which damaging climate-change events occur, it does place them further in the (physical) future than they would

²⁷This is of course not to say that no plausible theory of value can judge that mere (emissions-era) population spreading would result in an improvement, since there are all sorts of reasons for thinking that such population spreading would in practice affect lifetime well-being levels, despite the fact that (by stipulation) cumulative emissions are unaffected. I turn to some of these reasons in section 6.

otherwise be. Therefore, by pure discounting, it makes their negative impact on overall value smaller.

As it happens (I claim), the assumption of pure discounting is false. For example, as predicted above (since this theory ‘assigns intrinsic importance to physical time’), it too violates the Pareto Principle, although not quite as flagrantly as Instantaneous Averagism (Cowen, 1992),(Greaves, 2017, sec. 7.2). The plausibility or otherwise of pure discounting is a matter of well-worn debate, which I will not delve into in further detail here.²⁸ But we need not rely on that somewhat controversial claim here, because the argument I have sketched fails even if the assumption of pure discounting is granted.

The point is that if we seek to avoid the problems discussed in section 5.1, we are presumably starting from something like a discounted *totalist* value function, rather than a discounted instantaneous-averagist one. According to discounted totalism, the value of a state of affairs is given by the time-integral of the *total* amount of discounted instantaneous well-being at each instant of time. (More formally, the value function for discounted totalism is $\sum_t \sum_{i \in P(x,t)} \Delta(t) \cdot w_i^{inst}(x, t)$, where $\Delta(t)$ is the ‘discount factor’ for well-being at time t . Ordinary (undiscounted) totalism corresponds to the case in which $\Delta(t) = 1$ for all t ; otherwise, it is normally assumed that Δ is a decreasing function of t .) But if this is the assumed value function, it becomes crucial that emissions-era population spreading shifts *positive* contributions to well-being further into the physical future, to just the same extent that it shifts climate impacts futurewards. Because of this, a discounted totalist axiology is actually likely to prefer scenario A to scenario B: provided only that expected average per-capita well-being at each time remains positive, it is better by the lights of discounted totalism for more of it to happen sooner rather than later (in physical time).

²⁸Here I agree with almost all moral philosophers who have written on the topic of discounting (e.g. (Sidgwick, 1890; Cowen & Parfit, 1992; Broome, 2008), and with a healthy proportion of eminent economists (Ramsey, 1928; Pigou, 1932; Harrod, 1948; Solow, 1974; Cline, 1992; Dasgupta, 2008; Dietz *et al.*, 2008; Gollier, 2013), although those who dissent are also in distinguished company (e.g. (Arrow, 1999)). For a survey this debate, see (Greaves, 2017).

As is well known, taking account of exogenous risks of extinction in an expected-value framework leads to an effect structurally identical to positive pure discounting. This is a relatively uncontroversial contribution to the discount rate in practice, and the discussion below can be interpreted purely in terms of the extinction-related discount rate.

In the present context, however, we should note that it is unclear why the pure discount rate (whether primitive or extinction-risk-related) that we hold fixed across the scenarios of interest should be a rate relative to physical time, as opposed to relative to population time. If it is the discount rate relative to population time that is held fixed, a discounted version of totalism will agree with the undiscounted version that ‘spreading people out in time’, holding fixed individuals’ well-being levels, amounts to neither an improvement nor a disimprovement.

6 Instrumental reasons for favouring a physical-time slowdown

Section 5 surveyed and dismissed some possible reasons for thinking that *even if the assumptions of the simple model set out in section 4.2 were correct*, population reduction would lead to an improvement in the state of affairs as far as considerations of climate change are concerned (i.e. that population reduction would, in that axiological sense, ‘help with climate change’). I argued that both of the reasons suggested must be rejected.

The more promising reasons for thinking that population reduction might help with climate change take the form of reasons for doubting the assumptions of the model. In this section, I will examine each assumption in turn. In particular, I will discuss the case for thinking that the first and the third assumptions both fail. I also flag (in section 6.4) one further way in which population spreading could affect overall value, via its effects on future instantaneous population sizes.

6.1 (MPS-EE) is false: Effects of emissions-era population spreading on cumulative emissions

Recall the simple model’s first assumption:

(MPS-EE): Reducing instantaneous population size would implement mere emissions-era population spreading.

Recall further that (given my definition of ‘mere emissions-era population spreading’) it follows, from this assumption, that reducing instantaneous population size would not change the lifetime carbon emissions of the n th person, for any n .

This condition, and hence the assumption that entails it, is false. There are two reasons why, in reality, reducing instantaneous population size would be likely to lead to lower per-capita carbon emissions.

First: suppose we hold fixed the assignment of *energy consumption* rates to ordinally identified persons. (Suppose, that is, that any development of less energy-intensive technologies and lifestyles proceeds at a fixed rate in population time. We will revisit this assumption shortly.) A given person’s energy consumption must be met from a (possibly degenerate) combination of fossil fuel and non-fossil sources. But many of the non-fossil energy sources – for example, solar, wind and tidal power – are naturally constrained by physical time.

The amount of solar energy that can be captured by existing solar panel stocks, for example, proceeds at a fixed rate in physical time, paying no attention whatsoever to the number of humans who happen to be alive at that or any other time.²⁹ The upshot of this is that if we practice population spreading during the period before the achievement of zero carbon emissions, then a higher portion of the average person's energy consumption can be met from non-carbon sources – thus cumulative carbon emissions, and hence peak warming, can be reduced.

The second relevant consideration concerns the pace of research.³⁰ To get to a zero-carbon economy, we need research and development: assuming that a return to pre-industrial lifestyles would be unrealistic even if it were desirable, we need new technologies, to improve energy efficiencies and to decarbonise energy sources (including building up capacity to harness existing renewable energy sources), and thus to render something like modern life consistent with zero net carbon emissions. Since such research work is carried out by people, and a smaller instantaneous population would inevitably mean fewer researchers working at any given time, research will inevitably progress more slowly *in physical time* if near-term instantaneous population is reduced. This observation, though, is of limited direct interest; the more relevant question is what the effect of population spreading would be on the pace of research progress *in population time*. If population spreading would speed up the relevant research and development *relative to population time*, then this would provide a second mechanism that tends to reduce the emissions of a given ordinally identified person, and thus tends to reduce cumulative emissions.

This more relevant question is also substantially more complicated. On the one hand, there are some benefits to research projects and sub-projects being carried out in series rather than in parallel, so that later (sub-)projects can learn from and utilise the successes and failures of earlier ones; this consideration tends to lead to faster (population-time) research progress in scenarios involving lower instantaneous populations. On the other hand, though, the proportion of the population employed in research and development is presumably an increasing function of instantaneous population size, due to economies of scale elsewhere in the economy, and this consideration points in the opposite direction. (To take an extreme illustrative case: a world population of 10 would have its work cut out merely producing enough food for self-sustenance, and would be unlikely to support any research into the efficiency of solar panels.)

²⁹Of course, solar panel stocks (and their efficiency) can be increased, and, since such improvements are generated by people, rates of increase depend somewhat on population size. This is an important dependency; I return to it in the next paragraph. But (i) at some point the natural limits, and hence the physical-time flow rate of available energy, become relevant, and (ii) the *current* level of technology and capital is fixed, so that at least for the near term the flow of *harnessable* renewable energy is approximately independent of feasible variations in instantaneous population size.

³⁰Here, I discuss research in the service of mitigation. Similar considerations apply to research in the service of adaptation; cf. section 6.3.

We have, then, one set of considerations favouring the hypothesis that near-term population spreading would reduce cumulative emissions, and a second set whose sign is unclear. It is also unclear what the sign is of the net effect when both sets of considerations are combined. My own guess is that the effect of increasing the quota of renewable energy per person is greater than the net effect of population reduction on population-time research progress, so that one should indeed expect population reduction to reduce cumulative emissions. But this guess is based neither on hard data nor on rigorous theory, and there is scope for reasonable disagreement.

6.2 (CE-PH): Cumulative emissions and peak warming

The simple model's second assumption was that

(CE-PH) Peak warming is determined by cumulative emissions alone.

For present purposes, two limitations of the scientific case for this claim should be flagged.

Firstly, the models employed by the literature in question, and correspondingly the conclusion (CE-PH), apply only to emissions pathways for *carbon dioxide*, not to emissions pathways for GHGs in full generality. This is of some importance in the context of my discussion: emissions rates may well have more direct significance for shorter-lived GHGs, notably methane³¹. The reason for focussing nonetheless on carbon dioxide is that in practice CO₂ emissions constitute the majority of GHG emissions by a significant margin, on a wide variety of metrics, so that conclusions we reach on the basis of considering carbon alone are likely to be reasonably close approximations to the full story. In particular, Bowerman (2013, esp. fig. 7.5) shows that in e.g. the RCP4.5 emissions scenario³², for emissions in each of the next few decades, the contribution of CO₂

³¹For example: since carbon dioxide is a particularly long-lived greenhouse gas, making moderate reductions to the emissions rate *while preserving cumulative emissions* would not significantly affect peak CO₂ concentrations. But the same cannot be said for methane. It is not immediately obvious what significance (if any) such variables as peak concentrations have: this issue depends on the details of the climate system, and would be most naturally dealt with by a fresh modelling exercise. But since the relationship between atmospheric GHG concentrations and warming effects is not in general linear, they could easily be significant.

³²Much of the literature on climate change, including the IPCC's Fifth Assessment Report, focusses analysis on a standardised set of emissions scenarios ('representative concentration pathways', or RCPs): RCP2.6, RCP4.5, RCP6.0 and RCP8.0. The number following the letters 'RCP' is the level of radiative forcing in 2100 that is postulated by the scenario in question (in Wm^{-2} , relative to pre-industrial times); thus lower numbers correspond to more mitigation, and higher numbers to more warming. The scenario RCP4.5 has emissions peak around 2040, and is a fairly 'middle of the road' scenario.

emissions to peak warming is around four times the contribution from emissions of all other GHGs combined.

Secondly, even restricting attention to carbon emissions alone, the evidence for (CE-PH) consists in considering *relatively modest* perturbations to a chosen ‘status quo’ emissions pathway (for fixed cumulative emissions), and noting that when climate models are run with these perturbed emissions pathways as inputs, the output temperature pathway exhibits essentially unchanged peak warming. It is an open question what the limits of validity of the claim (CE-PH) are, when more radical variations in the emissions pathway (still holding fixed cumulative emissions) are considered. Heuristically, there are some physical-scientific reasons for suspecting that sufficiently radical variations in emissions timing would in fact break the link between cumulative emissions and peak warming (stemming from the fact within a few decades of CO₂ emissions, a significant proportion of the emitted CO₂ is absorbed by the land and oceans ((IPCC, 2013, Box 6.1, pp.472-3). However, since realistically achievable variations in the population pathway are also relatively modest (for example, the differences between the UN high/medium/low population projections), no achievable amount of population reduction is likely to take us outside the domain of validity of the claim (CE-PH).

6.3 (PH-CI) is false: Effects of population spreading on adaptation

Even if (contra the tentative conclusion of section 6.1) population spreading has no effect on *cumulative* emissions and hence (according to (CE-PH)) on *peak* warming, it is all but certain to reduce the (physical-time) emissions *rate*, and hence the (physical-time) *rate* of warming. This could be important, because some of the routes by which climate change is expected to lead to impacts on well-being are sensitive, in particular, specifically to the *physical-time* rate of temperature change.

Consider, for instance, species migration. One of the expected damaging effects of climate change is species extinction. And one prominent expected route to species extinction is the geographical migration of climate, proceeding at a *pace* that outstrips that at which the species itself is able to migrate, so that species are unable to keep up with the migration of what would otherwise be their viable habitat. Thus, for example, tree species, rodents, freshwater molluscs and carnivorous mammals can migrate at speeds up to about (respectively) 1km/decade, 12km/decade, 30km/decade and 60 km/decade. Meanwhile, projected climate velocities (depending on type of habitat, as well as emissions scenario) can easily be as high as 10-20km/decade, and could possibly be as high as 70km/decade, so that there is a real threat of some of these species being rendered extinct simply by their inability to keep up with the physical

pace of climate migration (IPCC, 2014a, p.15). The point here is that the fixed maximum speed at which a species is able to migrate is a speed relative to physical time, not to population time. (If rodents are able to migrate at a speed of 12km per physical decade, that would generally continue to be the case in a scenario in which the instantaneous human population doubled or halved; for instance, rodents would not suddenly develop an ability to migrate at 24km per physical year just because twice as many humans were on the planet.)

For another example, take agricultural outputs. Experts estimate that temperature increases of 2°C or more would be likely significantly to decrease yields of several major food crops, such as wheat, rice and maize (IPCC, 2014a, p.17). Adaptation to this impact might consist in genetically engineering higher-yield crops to offset or reverse the damage, and/or in switching to new foodstuffs that are more suited to growing in the new climate regime. Although this is less clear than in the case of species migration, it is also plausible that these adaptation processes, too, are at least somewhat constrained by physical as opposed to population time (for instance, because one round of experimentation with a new crop takes one physical growing season, independently of population size).

The general picture is that (1) to at least some significant extent, what is damaging about climate change is that it is a *change* we would need to adapt to, rather than necessarily that a resulting warmer climate would *in the long run* be forever worse than a cooler one. Meanwhile, (2) adaptation processes are often limited by rates in physical time in addition to (or, as in the case of species migration, instead of) being limited by rates in population time. To the extent that these two things are true, population spreading would be likely to reduce climate impacts even if it left cumulative emissions, and thus peak warming, unaffected. The simple model's assumption (PH-CI) is false.

It is perhaps worth noting that this latter explanation of how population spreading might reduce climate-change impacts is one that would not be captured by standard integrated assessment models of climate change, in their present forms. This is because the explanation essentially depends on the claim that the magnitude of climate damage at a given time t is positively related to the *time-derivative* of $\Delta T(t)$, not only to $\Delta T(t)$ itself (where $\Delta T(t)$ is the temperature rise at t , relative to some specified baseline date t_0 (e.g., 1990)). In contrast, the standard damage functions in both the MERGE model (Manne & Richels, 2004) and the DICE models ((Nordhaus, 2008, p.205),(Nordhaus & Sztorc, 2013, p.10)), for instance, hypothesise that climate damages at t are simply a quadratic function of $\Delta T(t)$.

6.4 Effects of population spreading on future population size

Another relevant consideration is that in practice, changes to the instantaneous population size during the emissions era would inevitably bring about changes to the population size during the times at which climate impacts are suffered. This is another route by which population spreading could affect the magnitude of climate damages, in evaluative terms, even without affecting impacts at a more purely physical level of description (e.g., the number and physical severity of floods). A given flood, for instance, causes more damage if there are more people around to be affected by it.

This at first sight might seem to suggest a further argument for the claim that population reduction would generate a climate-change-related improvement. The argument would run as follows: population reduction in the short to medium term would in practice result in lower population sizes during climate impacts; hence (for fixed climate impacts) it would lead to lower climate damages (in the sense of: lower climate-induced reductions in total well-being); hence it would generate a climate-change-related improvement.

This argument, though, is too quick. The comparisons it alludes to involve not two, but *four* states of affairs: high and low population scenarios, with and without climate change. The observation in this argument is that the value difference between the ‘high population, no climate change’ scenario and the ‘high population, climate change’ scenario is greater than the value difference between the ‘low population, no climate change’ scenario and the ‘low population, climate change’ scenario, when climate impacts are held fixed. And *this* is correct, insofar as ‘climate impacts’ are effects on average well-being. In that sense, reducing population (while holding fixed climate impacts) would reduce ‘climate damages’. But it does not follow that the ‘high population, climate change’ scenario is worse than the ‘low population, climate change’ scenario. The latter comparison is the one of interest for practical purposes, since of course we cannot turn off climate change.

Would it in fact be better to have smaller population sizes in the future, when climate impacts are suffered? This question is of a piece with the more general question of optimum instantaneous population size, and has nothing in particular to do with climate change. It therefore lies outside the scope of this paper.

7 Conclusions

A common line of thought identifies increasing human population size as one of the key drivers of climate change. The argument for this appears at first sight to

be simple: climate change is driven by GHG emissions, and, when average per capita emissions rates are held fixed, the global GHG emission rate is proportional to population size, so that at least in one clear sense, higher population means ‘higher emissions’. In this paper I have argued that while this is true, *that* sense of ‘higher emissions’ is of limited direct significance for the purpose of climate damages, so that the simple argument is misguided. Given the close physical relationship between degrees of warming and cumulative emissions, a more illuminating approach conceptualises increases and decreases in instantaneous population size in terms of speeding up or slowing down (what I have called) ‘population time’, relative to ordinary physical time. This conceptualisation makes clear that the key questions concern (1) the extent to which ‘emissions-era population spreading’ – that is, reducing instantaneous population size, and thus spreading out the people who will live before the goal of ‘zero net carbon emissions’ is attained more thinly over physical time – would in practice reduce cumulative emissions, and (2) the extent to which climate impacts at a given time are functions of the time-derivative of temperature, rather than only of the number of degrees of temperature rise itself. There are in the end credible reasons, from each of these two directions, for thinking that population reduction would indeed reduce climate-change impacts. But they are not reasons that are visible in the simple argument; they are significantly more subtle (and less conclusive) than the simple argument recognises.

Even if it is in the end correct that population reduction would help with climate change, this is of course not to say that a lower instantaneous population in the medium term would be better *all things considered*, because climate change is of course not the only relevant consideration for this general question. In particular, in this paper I have not considered:

1. The effects of instantaneous population size on average instantaneous well-being. Here there are two competing sets of considerations. On the one hand, economies of scale tend to favour larger populations. This is the dominant consideration at sufficiently low population sizes. On the other, beyond a certain population size, the more relevant consideration concerns fixed factors that cannot simply be scaled up to keep pace with population (for example, land area), and that eventually generate *diseconomies* of scale, together with the possibility of e.g. conflict over increasingly scarce resources. It is not clear (even approximately) which instantaneous population size would maximise average instantaneous well-being (in practice, given current conditions). There are also dynamic considerations, stemming from the effects of population *growth rates* (as opposed to population *sizes*) on dependency ratios.
2. Non-climate routes by which instantaneous population size may affect environmental degradation, and thereby affect future population sizes and/or well-being levels. Relevant considerations here include biodiver-

sity, air pollution, soil fertility, deforestation, groundwater, and fuel and mineral supplies.

3. The effects of instantaneous population size on the timeless population size. From the point of view of many theories of population axiology — specifically, theories that favour larger over smaller timeless populations when average well-being levels are held fixed — these effects are potentially very important to the general question of which instantaneous population size is optimal. Again there are two competing sets of considerations. On the one hand, insofar as the time profile of extinction risks for humanity is independent of near-term instantaneous population size, there is a *ceteris paribus* advantage to larger instantaneous populations: under that condition, larger instantaneous populations correspond to larger expected timeless population sizes. On the other hand, there are various mechanisms via which larger instantaneous populations in the short to medium term might increase anthropogenic extinction risks, and might thus decrease the expected size of the timeless population.

It is of course increasingly fashionable to claim that the world is ‘overpopulated’ in an all-things-considered sense, climate change merely providing one reason among a highly overdetermined case for thinking so. And this *may* be correct. But such claims are rarely based on sober analysis, and still more rarely are they sensitive to the highly relevant vagaries of population axiology (in particular, to the sense in which many axiologies assign intrinsic disvalue to population reduction). These are extremely complex matters, which I have not attempted to settle here. The aim of this paper has been only to clarify one piece of the jigsaw.

Acknowledgements

For helpful comments and questions, I am grateful to Ralf Bader, Niel Bowerman, John Broome, Mark Budolfson, Owen Cotton-Barratt and Teru Thomas, and to the audience of a lecture delivered at the University of Cambridge.

References

- Allen, Myles R, Frame, David J, Huntingford, Chris, Jones, Chris D, Lowe, Jason A, Meinshausen, Malte, & Meinshausen, Nicolai. 2009. Warming caused by cumulative carbon emissions towards the trillionth tonne. *Nature*, **458**(7242), 1163–1166.

- Archer, David. 2005. Fate of fossil fuel CO₂ in geologic time. *Journal of geophysical research: Oceans*, **110**(C9).
- Arrhenius, Gustaf. *Population ethics: The challenge of future generations*. Unpublished manuscript.
- Arrhenius, Gustaf. 2000. An impossibility theorem for welfarist axiologies. *Economics and philosophy*, **16**(2), 247–266.
- Arrow, Kenneth J. 1999. Discounting, morality, and gaming. In: Portney, Paul R., & Weyant, John P. (eds), *Discounting and intergenerational equity*. Washington: RFF Press.
- Bowerman, Niel. 2013. *Emissions targets for avoiding dangerous climate change*. Ph.D. thesis, University of Oxford.
- Broome, J. 2008. Can there be a preference-based utilitarianism? Pages 221–238 of: Fleurbaey, Marc, Salles, Maurice, & Weymark, John (eds), *Justice, political liberalism and utilitarianism: Themes from Harsanyi and Rawls*.
- Broome, John. 1991. *Weighing goods*. Oxford: Blackwell.
- Carlson, Erik. 1998. Mere addition and two trilemmas of population ethics. *Economics and philosophy*, **14**(2), 283–306.
- Cline, William R. 1992. *The economics of global warming*. Washington, D.C.: Institute for International Economics.
- Cowen, Tyler. 1992. Consequentialism implies a zero rate of intergenerational discount. In: Laslett, Peter, & Fishkin, James S. (eds), *Justice between age groups and generations*. Yale University Press.
- Cowen, Tyler, & Parfit, Derek. 1992. Against the social discount rate. In: Laslett, Peter, & Fishkin, James S. (eds), *Justice between age groups and generations*. Yale University Press.
- Dasgupta, Partha. 2008. Discounting climate change. *Journal of risk and uncertainty*, **37**(December), 141–169.
- Dietz, Simon, Hepburn, Cameron, & Stern, Nicholas. 2008. Economics, ethics and climate change. In: Basu, Kaushik, & Kanbur, Ravi (eds), *Arguments for a better world: Essays in honour of Amartya Sen. volume II: Society, institutions and development*. Oxford: Oxford University Press.
- Eby, M, Zickfeld, K, Montenegro, A, Archer, D, Meissner, KJ, & Weaver, AJ. 2009. Lifetime of anthropogenic climate change: millennial time scales of potential co₂ and surface temperature perturbations. *Journal of climate*, **22**(10), 2501–2511.
- Gollier, Christian. 2013. *Pricing the planet's future: The economics of discounting in an uncertain world*. Princeton University Press.

- Greaves, Hilary. 2017. Discounting for public policy: A survey. *Economics and philosophy*. Published online 31 May 2017.
- Greaves, Hilary. 2018. Optimal population size. *In: Arrhenius, Gustaf, Bykvist, Krister, & Campbell, Tim (eds), Oxford handbook of population ethics*. Oxford University Press. Forthcoming 2018.
- Harrod, Roy. 1948. *Towards a dynamic economics*. London: Macmillan.
- Hurka, T. M. 1982a. Average utilitarianisms. *Analysis*, **42**(2), 65–69.
- Hurka, T. M. 1982b. More average utilitarianisms. *Analysis*, **42**(3), 115–119.
- IPCC. 2013. *Climate change 2013: The physical science basis. Working Group I contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge, UK and New York, NY, USA: Cambridge University Press.
- IPCC. 2014a. *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*. Cambridge, UK and New York, NY, USA: Cambridge University Press.
- IPCC. 2014b. *Climate change 2014: Mitigation of climate change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press., Cambridge, UK and New York, NY, USA.
- Manne, Alan S., & Richels, Richard G. 2004 (June). *MERGE: An integrated assessment model for global climate change*. Available online at <<http://web.stanford.edu/group/MERGE/GERAD1.pdf>>. Accessed 30 June 2017.
- Matthews, H Damon, & Caldeira, Ken. 2008. Stabilizing climate requires near-zero emissions. *Geophysical research letters*, **35**(4).
- Matthews, H Damon, Gillett, Nathan P, Stott, Peter A, & Zickfeld, Kirsten. 2009. The proportionality of global warming to cumulative carbon emissions. *Nature*, **459**(7248), 829–832.
- Nations, United. 2017. *World population prospects: The 2017 revision. key findings and advance tables*. Tech. rept. United Nations, Department of Economic and Social Affairs, Population Division. Working paper No. ESA/P/WP/248.
- Ng, Yew-kwang. 1989. What Should We Do About Future Generations? *Economics and philosophy*, **5**(2), 235–253.
- Nordhaus, William. 2008. *A question of balance*. New Haven, CT: Yale University Press.

- Nordhaus, William, & Sutorc, Paul. 2013 (October). *DICE2013R: Introduction and user's manual*. Available online from <http://www.econ.yale.edu/nordhaus/homepage/documents/DICE_Manual_100413r1.pdf>. Accessed 30 June 2017.
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Clarendon Press.
- Pigou, Arthur C. 1932. *The economics of welfare: Volume 1*. London: Macmillan. 4th edition.
- Population Matters. 2016 (May). *Population policy and the environment: Joint international position statement, version 10*. Available online from <https://www.populationmatters.org/documents/position_statement.pdf>. Accessed 30 June 2017.
- Porrirt, Jonathan. 2010. Population and climate change. *New internationalist*, **429**(January).
- Ramsey, Frank P. 1928. A mathematical theory of saving. *Economic journal*, **38**, 543–559. Reprinted in F. P. Ramsey, *Foundations: Essays in philosophy, logic, mathematics, and economics*, ed. D. H. Mellor.
- Royal Society. 2012. *People and the planet*. Tech. rept. Royal Society Science Policy Centre. Report 01/12.
- Sidgwick, Henry. 1890. *The methods of ethics*. London: Macmillan.
- Solow, Robert. 1974. The economics of resources or the resources of economics. *American economic review papers and proceedings*, **64**, 1–14.
- United Nations. 2015. *World population prospects: The 2015 revision. volume I: Comprehensive tables (ST/ESA/SER.A/379)*. Tech. rept. United Nations, Department of Economic and Social Affairs, Population Division.
- Zickfeld, Kirsten, Eby, Michael, Matthews, H Damon, & Weaver, Andrew J. 2009. Setting cumulative emissions targets to reduce the risk of dangerous climate change. *Proceedings of the national academy of sciences*, **106**(38), 16129–16134.