

Existential Risk, Astronomical Waste, and the Reasonableness of a Pure Time Preference for Wellbeing

SJ Beard and Patrick Kaczmarek, Centre for the Study of Existential Risk, University of Cambridge

In this paper, we argue that our moral concern for future wellbeing should reduce over time due to important practical considerations about how humans interact with spacetime. After surveying several of these considerations (around equality, special duties, existential contingency, and overlapping moral concern) we develop a set of core principles that can both explain their moral significance and highlight why this is inherently bound up with our relationship with spacetime. These relate to the equitable distribution of:

1. Moral concern in a universe that we can only causally affect in one temporal direction;
2. Access to the benefits from using spacetime as a resource; and
3. The burdens of care given 1 and 2

We conclude by considering the practical implications of our argument and find that, while it is often assumed that a preference for present over future wellbeing weakens the case for existential risk mitigation, this likely is not the case.

“Make me this present then: your hand in mine/and we'll live out our lives in it.”
- from *The Present* by Michael Donaghy

1 | Introduction

Imagine that you are entirely indifferent about when you benefit from consuming the resources available to you and that you live in a world where resources you leave unconsumed will generally increase over time (plants and animals grow and reproduce, share prices rise, investments in research and development lead to new technologies, and so forth). Under idealized circumstances, your indifference to when you reap the benefits of these gains would mean that you would always decide to invest as much as you could in the future and consume as little as possible now. But while doing so would lead to the maximum possible level of available resources, it would push the time of their consumption indefinitely far into the future. Like a miser you would grow fabulously rich while living in abject poverty.

Now, instead, imagine that you strongly prefer to consume your resources right now, so much so that the prospective gains of future growth mean nothing to you. You would then decide to consume as much as possible and invest as little as you could (likely a negative amount, taking on debts to fund your insatiable splurge). Once again, you would be left in abject poverty, albeit this time due to imprudence rather than excessive prudence.

The optimum, in the sense of wellbeing maximising, strategy to pursue under such circumstances is to avoid these extremes and to save something, but not everything. More specifically, the optimum strategy is to save so that the rate of your saving multiplied by the marginal utility of your consumption is always equal to the difference between your current enjoyment of the goods you consume and the maximum possible level of enjoyment you could have (i.e., the level of enjoyment you would experience when your declining marginal utility of consumption reaches zero). Thus, as your savings increase you should choose to save less and less of your overall income, until you reach the point at which consuming

goods ceases to make you any happier. At that point you have no need to save any more for the future and can enjoy a never-ending state of “bliss” (Dasgupta 2019a).

The above paragraphs offer a highly simplified account of the basic intuition behind Frank Ramsey’s 1928 paper, *A Mathematical Theory of Saving*, which is recognized by many as the origin of the contemporary economic approach to ‘discounting’. While not given in the paper itself, the appropriate attitude to saving that Ramsey sets out in that paper would be governed by the following social discount rate for future consumption:

$$r = \delta + \eta g$$

where g is the expected rate of growth and the so-called ‘taste parameters’, η and δ , represent the (declining) marginal wellbeing of consumption and our pure time preference for wellbeing respectively.

Ramsey states that, in an ideal world, δ should be equal to 0 and that a pure time preference for wellbeing is “a practice which is ethically indefensible and arises merely from the weakness of the imagination” (1928, 543).¹ Nevertheless, he included δ in his formulas for two reasons. Firstly, because if an individual had such a pure time preference for wellbeing, then the optimal strategy *for them* would need to take account of this. Secondly, and more importantly, because if we apply optimal savings theory, and by extension discounting, to social choices, then we must take account of people’s actual preferences. In particular, the long-term interest rate in a society will be determined by the discount rate of the “marginal saver”; that is the person who, at this level of consumption and economic growth, is indifferent between saving or borrowing. As such, the social discount rate must take account of this person’s pure time preference for wellbeing, whether that is ethically defensible or not.

This is known as the ‘positivist’ approach to determining the pure time preference for wellbeing, and while Ramsey’s formulas for optimal saving have met with near universal admiration and agreement, this element has been the source of considerable controversy. Some, mainly moral philosophers, have argued that even respecting people’s observed time preference for wellbeing is immoral and that δ ought to be set to 0 when we are applying a ‘social’ discount rate, for the purposes of public policy rather than private choices. Others, mainly economists, believe that Ramsey was mistaken in claiming the pure time preference was ‘ethically indefensible’, and that there are good moral reasons why people have such a preference and why this should be respected; this is known as the ‘normativist’ approach to discounting. While philosophers have tended only to criticise the normativist position, we believe this criticism has often been far too hasty and that moral philosophers have

¹ Intriguingly, Ramsey’s personal views may not have reflected this much cited claim. His statements merely repeat the standing economic orthodoxy, derived by Arthur Pigou from utilitarian principles, that discounting was due to our “telescopic faculty” and that “the State should protect the interests of the future in some degree against the effects of our irrational discounting and of our preference for ourselves over our descendants” (Pigou 1932, 25/29). Yet, Ramsey, who was personally opposed to any kind of ethical absolutism, may well have adopted these principles simply to provide a coherent moral framework for his arguments. Elsewhere, he puts things very differently: “My picture of the world is drawn in perspective and not like a model to scale. The foreground is occupied by human beings and the stars are as small as three-penny bits. I don’t believe in astronomy, except as a complicated description of part of the course of human and possibly animal sensation. I apply my perspective not merely to space but also to time. In time the world will cool and everything will die; but that is still a long time off and its present value at compound discount is almost nothing” (Ramsey 1931, 291).

something to offer to economists of the normativist persuasion if we are willing to constructively engage with and develop their arguments.

2 | A brief survey of normative arguments for discounting

It is sometimes assumed that, to be justified in discounting the value of future wellbeing, we must believe that the temporal location of wellbeing is relevant to its value in absolute terms, that present wellbeing is simply 'better' or more valuable than future wellbeing.² This would certainly strike us as an absurd position. However, to claim that it is the only reason one might be less concerned about future wellbeing is also absurd; we could discount the value of future wellbeing because the passage of time has intrinsic effects on us as temporally situated agents, and these effects then give us reasons to have differential concern for wellbeing that stands in a different temporal location relative to us (what Ramsey referred to as its 'modified utility').³

Compare utilitarianism, which in its usual forms is impartial. Very few utilitarians believe that a parent has absolutely no reason to care disproportionately about the wellbeing of their own children. But this is not because they believe their child's wellbeing is worth more than a stranger's wellbeing. Rather, a parent's special concern for their child's wellbeing may be grounded in reasons of practicality (parents can help their children in ways others can't), proximity (parents have many opportunities to be helpful to their children), the impersonal value of special relationships (if giving and receiving appropriate care is an aspect of the good life, and thus inherently part of wellbeing, e.g. (Parfit 2016, 123)), or other similar reasons. Nor does appealing to these reasons necessarily make one any less of a utilitarian; the moral facts they appeal to can all be articulated from an impartial perspective of concern for human wellbeing (even if they imply that individuals may exhibit that concern in different ways). We will argue that something similar is true in relation to the passage of time.

In the remainder of this section, we briefly consider four of the effects that the passage of time has upon us: (a) it creates opportunities for growth; (b) it imposes unbridgeable divisions between people; (c) it makes elements of the future existentially contingent upon events in the present; and (d) it restricts us to only one direction of causal influence. In the next section we attempt to build a more general argument about the ethical significance of causation to defend a more comprehensive account of why the passage of time is morally significant.

2.a. Discounting due to equality and growth

Many people are concerned with the distributive ideal of equality, the idea that everyone should be, or have, the same in some morally relevant respect.⁴ It has long been understood

² E.g., "the most common and straightforward argument for a zero pure time discount rate maintains simply that we have no reason to attribute fundamental moral importance to someone's location in time" (Caney 2014, 323).

³ See Kelleher (2017) for more on the ways in which economics and philosophers have tended to misunderstand each other on this point, Greaves (2017) for a more charitable philosophical engagement with the standard practise of discounting in economics, and Heath (2021, 254) for a different account of how our relationship with time may justify discounting.

⁴ Ramsey was certainly among them. He ended *A Mathematical Theory of Savings* by noting how differences in δ among different people would lead to inequality, because there would always be some who discount above

that such concerns might be relevant to our evaluation of future wellbeing. People in the future can be expected to be generally better off because of social, economic, and technological progress, re-generation, growth, and productivity in natural systems (from individual plants and animals to ecosystems like forests), and the asymmetry that each generation can improve the wellbeing of those that come after it but not of those who came before it. While it is not technically possible to tax the future to benefit the present, we can nevertheless take account of the moral significance of this inequality when weighing up the claims of people in the present and future. In fact, Ramsey himself drew attention to this point, noting that a pure time preference “has the effect of accelerating the decrease of marginal utility and lessening the relative importance of high income” and thus allows us to “translate our discounting of the future into a discounting of high incomes” (1928, 554).

The idea of discounting future wellbeing due to a concern for equality is common among economists. However, it is generally implemented by modifying η , the marginal utility of consumption, rather than δ , the pure time preference. For instance, William Nordhaus claimed that:

The second parameter that determines return on capital is the consumption elasticity, denoted as η . This parameter represents the aversion to the economic [inequality] among different generations. A low (high) value of η implies that decisions take little (much) heed about whether the future is richer or poorer than the present (2007, 202).

Building on this, he then suggested that a relatively high level for η should be used for discounting future costs and benefits.

Stern and Taylor in turn replied to this that, if that was so, then the value for η used by Nordhaus and other economists represented a level of egalitarian concern that was inconsistent with their other views:

In the case of $\eta = 3$ in Nordhaus' example, over 99% could be lost and a transfer [between a wealthy person and someone with only 20% of their resources] would still be beneficial. Does he advocate huge increases in transfers from rich to poor in the current generation? (2007, 203-204)

While agreeing that there is a link between discounting and egalitarian concerns, we believe this exchange highlights the problem with attempting to bake the ideals of equality or fairness into our discount rate via the value of η .

To recap from our introduction, discounting future consumption by ηg is optimal in the sense of producing savings behaviour that leads to the highest stream of wellbeing over time. However, an egalitarian concern is not about optimizing the value of utility but *modified utility*, the value of utility from the egalitarian perspective. The marginal utility of

the social rate, causing them to consume and borrow until they were reduced to a subsistence level of consumption, while others discount at or below the social rate, building up savings and thus increasing their standard of living. Yet, this was not meant to be his last word on the matter but was originally written to conclude part 1 of a longer paper that continued with his thoughts on how the state should raise and redistribute resources. Nevertheless, following the advice of Milton Keynes, Ramsey decided to split the papers, publishing his theory of taxation separately (1927) and noting that in *A Mathematical Theory of Savings* he would “ignore altogether distributional considerations” (1928, 543).

consumption is a value that it is possible to determine via empirical means and it should not be subject to the kind of debate Nordhaus and Stern engage with. Instead, we need to follow Ramsey's lead and translate our discounting of high incomes into discounting of the future by modifying the value of δ to account for expected growth.⁵

Not everyone is a fan of egalitarian based discounting, however. This includes Derek Parfit. In his influential appendix on the social discount rate, Parfit claims that, even if some people in the future will be better off than we are, others may be no better off, or even worse off, and it is wrong to discount these unfortunate people's wellbeing simply because they were born at the same time as others who fare better (1984, 484). However, an egalitarian should not only consider how well off those who may be benefited by our actions are, but also how well off those who might alternatively be burdened by our inaction would be as well. No generation can undo social injustices locked away in the past.⁶ But in a world in which the availability of resources is increasing, it is not unfair to expect those who benefit from this growth to be the ones who are given the principal responsibility for helping those who have not thus benefited. Perhaps many wealthy Victorians were better off than many who are now poor; however, had they made sacrifices for these currently poor people, then this would have been a heavy burden for them. Jeff Bezos, and other contemporary billionaires, on the other hand would still be extremely well off, by the standards of both the 21st century and the 19th, if they were to do much more to help the present-day poor. Yet, even these contemporary billionaires may not be as well off as many who will live in the future, and while fairness would seem to demand that they do more to tackle present day inequalities as a result of their current relative wealth, it may be unfair to expect contemporary billionaires to take on the additional burden of reducing future poverty as well when the riches that people in the future could choose to divert to this task may be so much greater.

2.b. Discounting due to agent relative special duties

As well as appealing to impersonal values, such as equality, we can also argue for discounting based on the kinds of relationships in which we stand to some people but not others. Special duties of this sort are old hat; for instance, most already believe that it is permissible to refuse to divert a runaway trolley into the path of one's own child to save five innocent strangers from a similar fate (Bleske-Rechek et al. 2010). What makes someone's choice to spare her own child permissible is not that the five strangers have a lesser moral status than her child in general but that, as a parent, she is in possession of agent-relative reasons concerning her own child. Indeed, according to at least some philosophers, turning the trolley into one's own child could be gravely wrong, even if doing so really was for the best (Kamm 2004, 674).

In general, we are bound by special obligations to those around us, not only to close friends and family members but also colleagues, co-citizens, people who we engage in

⁵ It is also worth noting that while the marginal utility of consumption is generally taken to be a continuous linear function, the moral value/modified utility of wellbeing may be quite different to this. A linear diminishing marginal value for wellbeing is generally known as prioritarianism. Other views include the moral value of wellbeing depending on the wellbeing of other people (egalitarianism), the moral value of wellbeing diminishing in a discontinuous fashion (sufficientarianism), and hybrid views (such as MaxiMin and LexiMin).

⁶ However, this is not to deny that we are able to do *something* about at least some past harms. See (Kaczmarek and Beard 2020).

collective projects with, and even simply those who we are in a unique position to aid.⁷ While we may have some obligations to people in the future, the passage of time has a general effect of weakening, or even severing, interpersonal relationships. Thus, such obligations are seldom as strong as our obligations to those we are currently involved with, and they will weaken over time.⁸

But, as Andreas Mogensen flags, while this helps explain why we care about descendants who are nearer to us in time, “it provides no justification for caring more about the welfare of unrelated strangers on the basis of their location in time” (2022, 2741). However, we are here interested in the social discount rate, not any individual’s discount rate, and Mogensen responds that, just as a parent is justified in caring more about their own child’s particular fate than the fates of strangers, the present generation may be justified in caring more about generations who are nearer to them in time. As Mogensen puts it, “We together may have greater reason to care about the next generation than about later generations, because those who are born into the next generation are our children, whereas succeeding generations will be more and more distantly related to those of us living now From the perspective of ‘the world community now,’ there are no strangers, present or future.” (2022, 2743/2745). Mogensen refers to this as ‘discounting for kinship’ and acknowledges that it is unlikely to serve as a full explanation for the pure time preference for wellbeing because “What matters ultimately is not distance in time, but in the family tree” (2022, 2745).

2.c. Discounting due to existential contingency

For projected wellbeing to be valuable, it must be realised in the right way, for instance, it must benefit some morally significant subject. The most obvious way in which this may not happen is if these benefits are never realised at all, because there are no people in the future to receive them. That we might care less about potential future wellbeing due to this risk is often taken to be the least controversial component of the social discount rate; for instance, even Stern incorporates an assumed background risk of human extinction of 0.1% per year⁹ as the sole component of δ (Stern 2007, 47).¹⁰

It is important to note that this rate is meant to represent the background risk of human extinction; that part of the risk that does not depend upon the choice under consideration. It can therefore be treated as an exogenous variable; unlike the risk of human extinction occurring *because of our actions*, which should not be used to discount

⁷ In his 2017 Uehiro lectures, Larry Temkin suggested a range of conditions for having special obligations that we might summarize under the label of ‘being in a unique position to aid’ (Temkin 2022). We build on the point in §3.

⁸ Arguably the most significant special relationship is with ourselves, and here again we may have stronger obligations to serve our present interests than our future interests.

⁹ Note that an extinction risk of 0.1% is far more significant than might at first be realized. For example, if this were the risk of human extinction today, then people would be more likely to die in a human extinction event than a car crash (Cotton-Barratt et al 2016).

¹⁰ Once more, it appears that Ramsey himself was cognisant of this reason for discounting. That we should discount future wellbeing according to the probability it would not materialize was a suggestion made to him by Keynes; however, Ramsey chose to neither accept nor argue against this point. Instead, he merely comments that “we have to assume that the community will always be governed by the same motives as regards accumulation, so that there is no chance of our savings being selfishly consumed by a subsequent generation; and that no misfortunes will occur to sweep away accumulations at any point in the relevant future” (1928, 544).

the value of future benefits but evaluated 'at cost', by considering all the wellbeing that would be forgone as a result (Ng 2016).

Despite being the most widely accepted reason for discounting future wellbeing, some philosophers refuse to go even this far in accepting a pure time preference, on the grounds that even the background risk of extinction should be incorporated directly into our assessment of an outcome's expected costs and benefits (Parfit 1984, 481-482). At one level, this is clearly no more than a terminological dispute as that is precisely what applying a pure time preference would achieve. However, we think including the background extinction risk into a pure time preference for wellbeing has two advantages over this proposal.

Firstly, the way that we need to account for this risk is highly dependent upon the existential contingency of the future on the present. What matters is the cumulative build-up of this risk over time, since humanity only needs to go extinct once to render the entire subsequent future worthless to us. An extinction risk at any point in time will thus devalue all the expected wellbeing after that point, even if we might otherwise be expected to eventually reach a state of total human safety. Hence, there is a very close relationship here between the passage of time and our reasons for discounting future wellbeing.

Secondly, it may be wrong to say that we only discount future wellbeing by the extent to which human extinction lowers the expected wellbeing of our choices. We may in fact not know the probability of this occurrence to any degree of certainty, and even if we can make some kind of assessment of the risk, we may prefer not to discount by this precise probability but rather to some higher degree (because we are risk or ambiguity averse) or to a lesser degree (because we care more about making the right choices with regards to those futures where humanity survives). Hence, we do not have to treat 1 unit of wellbeing now as of equal value to 2 units of wellbeing that could be expectedly realized with 50% probability in 100 years' time. Instead, we may have reasons to prefer the certainty of one unit of wellbeing now or to commit to some intergenerational project whose value will only be realized if humanity survives to complete it.

Another feature of existential contingency that may be relevant to our evaluation of future wellbeing relates not to the survival of humanity, but of individual human beings. There is a view that benefits to any not presently existing people do not matter.¹¹ This often gets summarised by the dictum that "We are in favour of making people happy, but neutral about making happy people" (Narveson 1973, 67). On the sternest of these views, it does not matter whether humanity survives from the decision-making perspective of those people now alive; once all the people who are currently living have died, future wellbeing will cease to have any value for us. If this was our view, then our pure time preference should not be sensitive to the background human extinction risk but instead to the background level of human mortality (approximately 0.8% per year in 2021).

More plausible perhaps is the view that we should attend to the wellbeing of possible people in the future, but to a lesser extent than that of those who are alive and living with us now, or that, among future people, we should care disproportionately more about making a given future person better or worse off than determining whether a better off possible person or a worse off possible person come into existence (Otsuka 2018).

¹¹ The most plausible version of this argument runs as follows: although future people matter, there is nothing we can do to help them since any action we take simply changes who comes into existence, which is not better for anyone.

Partha Dasgupta has argued for this position using the following case, which he calls 'Replacement' (Dasgupta 2019b, 65; cf. Hutchinson 2019). Imagine an individual who will soon die if we do not intervene. We can either save this person's life, in which case they will go on to live for 50 more happy years, or we can allow them to die and instead bring into existence another person who we are certain will have an equally happy life for the next 50 years. In terms of the quantity of wellbeing that will be created, these two choices are equivalent. However, many feel that we have stronger moral reason to save the life of the already living person, because doing so is not merely valuable but a clear benefit to them. On the other hand, given that the not currently existing person would shoulder no burden if they weren't brought into existence, since there would be no one "there" to bear that harm, their coming into existence is much less clearly beneficial for them.

Since the existential contingency of the future always increases over time, as the effect of these choices spreads out through the universe, this kind of view will lead us to generally prefer present wellbeing to future wellbeing, and thereby justify some form of pure time discounting. Again, however, this may not be uniform across lives, since the existence of some people (such as any children that we might have) is far more contingent upon our decisions than the existence of others (such as the children of strangers who live far away from us), but this might not matter for social, rather than personal, discounting.

2.d. Discounting due to causality and care

Harry Lloyd (2021) has recently suggested another justification for the pure time preference on the grounds that we have special duties to those with whom we share a so-called 'region of moral concern'.

A person's region of moral concern at location x and time t is the region of spacetime that it is *nomologically* possible for her to causally influence at time t , from location x (Lloyd 2021, 7). It describes her potential moral burden; what the moral agent should worry about when trying to work out the "longlist of events that she might be morally obligated at time t to influence" (Lloyd 2021, 8). Clearly, this person's location in time will dictate much of what falls outside of her region of moral concern. For example, I am under no obligation to thwart the Nazis' evil plans since time travel is impossible. However, one's location in space will also affect what falls in and out of this region. For example, if trapped in an underground bunker, shut off from the rest of the world, then I am incapable of rescuing a child from drowning, whether they are ten feet above me or a hundred years in the future. Similarly, since the effects of my actions can never spread out into the world faster than the speed of light, there will be some harms happening right now that I am unable to prevent, and thereby that I have no moral obligation to prevent, from happening.

Lloyd grounds the importance of this region in people sharing a *moral burden* in common (Lloyd 2021, 10). Plausibly, such shared burdens can bring people closer together, and in that respect, they resemble the grounds for special obligations that we discussed above.

On Lloyd's view, those currently alive will have a greater overlap in their regions of moral concern than they will with people in the future. He illustrates this with a toy example:

Bunker. Betty is trapped inside a bunker, where she will spend the rest of her life. Inside there are the controls for two missiles. The smaller missile is programmed to fire in five days' time, towards Town One, which has a population of 5,000 people.

The larger missile is programmed to fire in seventy years' time, towards Town Two, which will then have a population of 6,000 people. These missiles will destroy these towns unless deactivated. Betty can deactivate the smaller or larger missile, but not both.¹²

Figures 1, 2, and 3 describe the regions of spacetime it is nomologically possible to causally influence at time t , shaded grey, for Betty (R_B), the residents of Town One (R_{T1}), and future residents of Town Two (R_{T2}). These regions of concern cover an expanding area of space (represented by the x axis) over the passage of time (represented by the y axis) because of the limits on the speed at which we can affect things, as mentioned above, meaning these regions appear as triangles in the diagram.¹³

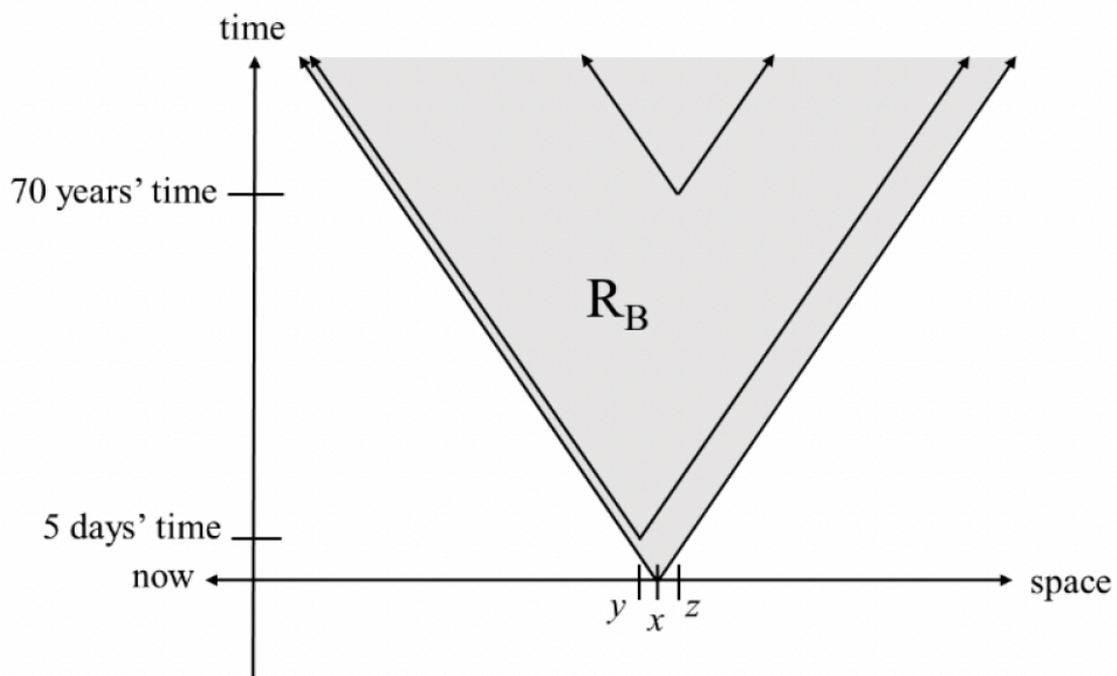


Figure 1. Source: Lloyd 2021, 9.

¹² This is a slightly modified version of the original (Lloyd 2021, 1).

¹³ Note that x , y , and z represent the spatial locations of Betty, Town One, and Town Two, respectively.

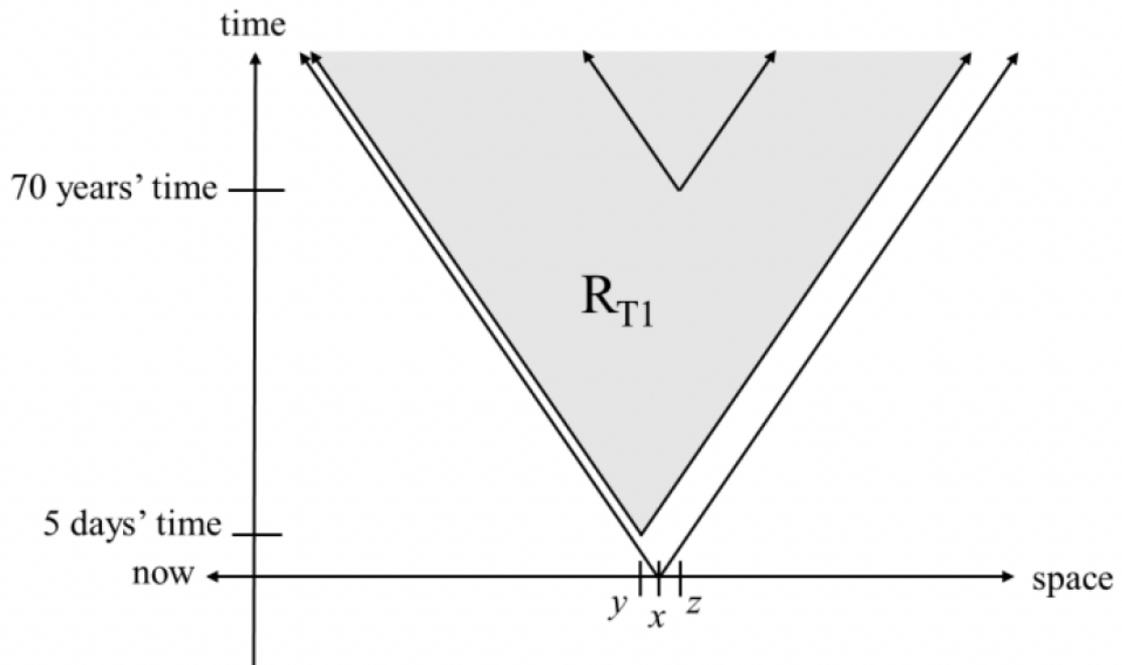


Figure 2. Source: Lloyd 2021, 9.

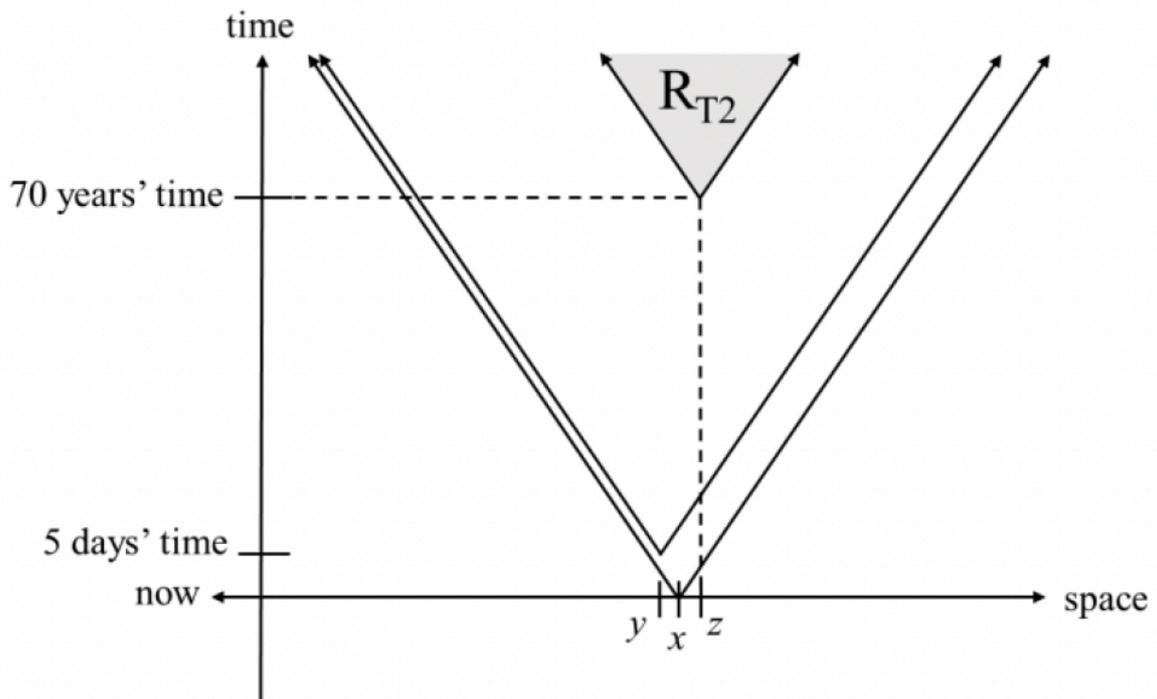


Figure 3. Source: Lloyd 2021, 9.

As can be seen in Figure 4, Betty has a greater region of moral concern in common with the residents of Town One than she does with the future residents of Town Two.

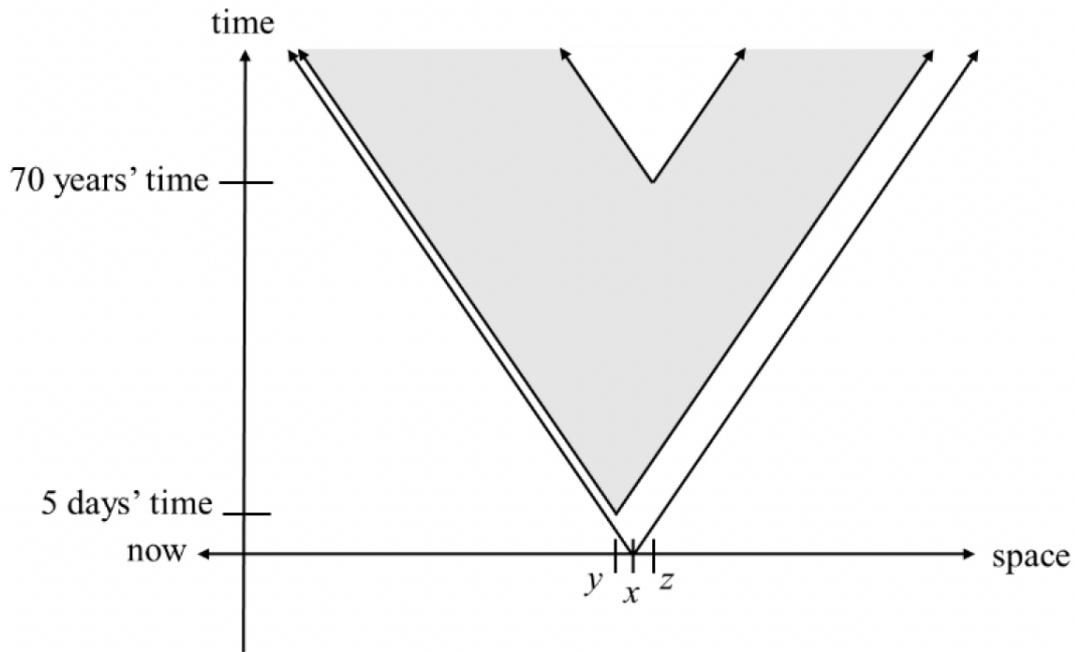


Figure 4. Source: Lloyd 2021, 10.

And so, on the view proposed by Lloyd, Betty has a stronger special duty to prevent the deaths of the residents of Town One than she does to save those people who will reside in Town Two.

While there is much to admire about Lloyd's proposal, we have reservations about its current formulation. The strength of this special duty depends upon the degree of overlap between the regions of spacetime whose events it is *nomologically possible* for people to have an obligation to influence. We cannot help but wonder why we should pick out nomological possibility, as opposed to some other kind of possibility, and furthermore why the strength of this special obligation does not instead depend on the overlap in actual burdens.

One reason Lloyd might want to avoid relying on regions of *actual* moral concern is to avoid the following kind of problem. Consider:

Bad Breakup. Ages ago, George and Jane divorced, swearing off any special duties between them as a couple; they are strangers to each other in both their own and morality's eyes. George has special duties to his children, Elroy, and Judy, as well as towards Astro (their children's dog) and Rosie (the maid). Jane has the same special duties towards Elroy, Judy, Astro, and Rosie as does George. Judy and Jane are trapped in a burning building, and George can save one or the other but not both.

In *Bad Breakup*, George shares a certain moral burden in common with his ex-wife, Jane, that he does not share with his daughter; namely, special duties towards his children, dog, and maid. And so, just as Betty has a stronger obligation to deactivate the smaller missile in Bunker, George's strongest obligation could be to rescue Jane from the flames. This conclusion strikes us as preposterous.¹⁴ George should clearly rescue Judy, whom he has a

¹⁴ In personal communication, Lloyd confirms that this sort of case was one of his motivations for focusing on nomological possibility.

strong prerogative to protect from harm, not the “stranger” who shares this special concern for his daughter.

We also have reservations about the intuitive plausibility of a shared region of moral concern strengthening (or grounding) any special obligations one has towards a fellow duty-bearer. Consider:

Slave Owners. Members of the present generation might have owned slaves. If they had, they would owe reparations to their ex-slaves. This would be a *pro tanto* duty that they do not share with future generations (since retrocausality is nomologically impossible).

We find it absurd to think that, in sheer virtue of potentially being partners in crime, any special consideration present people might legitimately owe each other should be amplified, and perhaps even be made strong enough to outweigh any actual debts to descendants of their ex-slaves. For this reason, a shared region of moral concern strikes us as a problematic source of special consideration.

Nevertheless, we do think there are at least some kinds of shared moral concern that ground special moral considerations, namely those burdens that are shared in virtue of working together to achieve common moral objectives or to complete ‘Burkean’ projects that are so large and significant that they require intergenerational co-operation (Beard and Kaczmarek forthcoming).

3 | Developing a generalized argument for discounting

The preceding section surveyed four kinds of arguments for why we should have a pure time preference for wellbeing, which we interpret as a greater degree of moral concern for the wellbeing of people right now compared to that which will exist in the future. They are often objected to for being insufficiently ‘pure’ in that these arguments rely on factors not inherent to the passage of time. As set out at the start of section 2, our view is that this may simply indicate a terminological dispute about what it means to have a pure time preference and we note that philosophers tend to be somewhat prudish in our reading of such ‘purity’ and economists less so.

However, we do agree that when taken individually there is a case to be made that the best response to each of these arguments might be to discount for something other than time (at least where we have a sufficient degree of information about the futures in question to do so). Yet, we do not have to read them all separately, and it is our view that these arguments are strongest when read together, not as competing accounts for why time supposedly matters in our moral deliberations. To our minds, they are best viewed as four ways in which the moral significance of time can be seen as manifesting. In this section we aim, not to repudiate and replace any of the above, but instead to build a general theory about the significance of time’s passing.

3.a. *The real issue about the distribution of care*

We will begin with Lloyd’s account of why we should discount future wellbeing, which we believe offers an invaluable jumping off point for understanding the fuller significance of time.

Lloyd notes that the way we experience spacetime only allows us to causally affect events that take place in our future, and that the physical limitations of our universe (such

as the speed of light) mean that the special area we can affect is also limited, although it expands into the future, as shown in Figure 5:

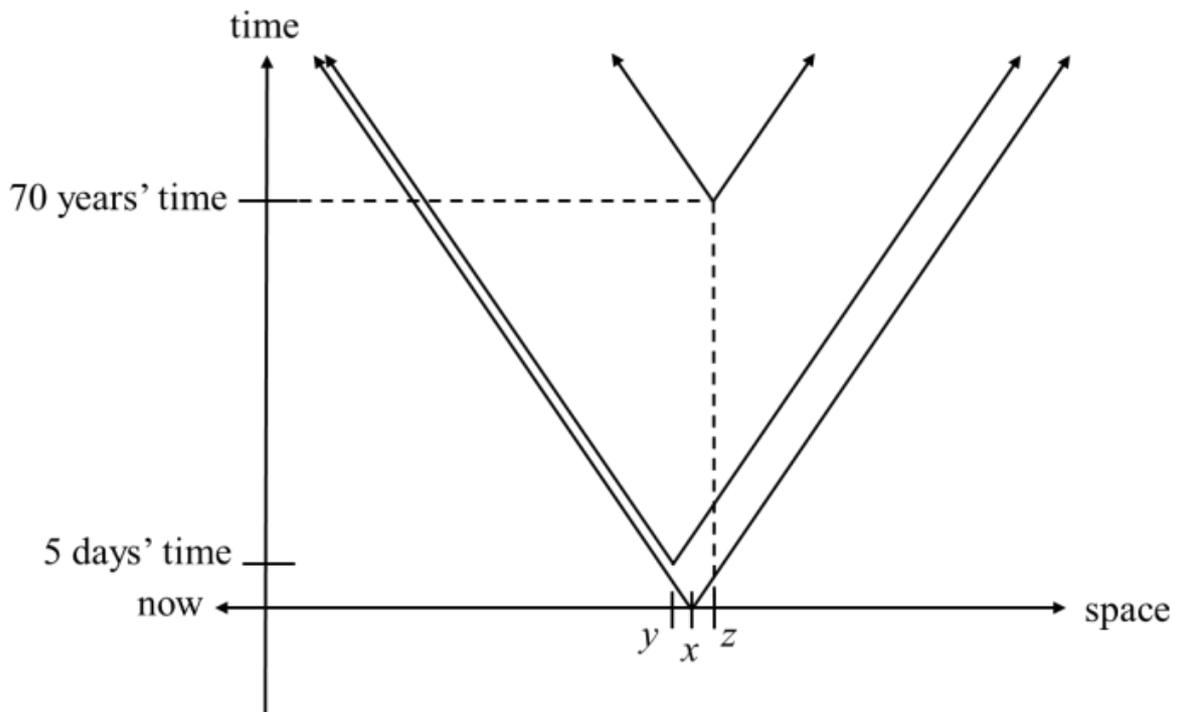


Figure 5. Source: Lloyd 2021, 8.

As described in the previous section, Lloyd builds his argument by considering the degree of overlap between three agents, whom we will call 'X', 'Y', and 'Z', given their locations in space (x , y , and z).

However, we believe that a more interesting feature of this graph is the size of these three areas. Even if all three agents were to live forever, the area that Y could causally influence will be fully contained within that which could be influenced by X, and that of Z will be entirely contained within both of these (since regions are defined by nomological possibility, if an agent is within another agent's region of concern, their entire region of concern will be as well). If we imagine a fourth agent W, who was sufficiently spatially distant from these three, then the area she could causally influence may not be fully contained within any of them (although they would eventually overlap). In that case X, Y, and Z may not be able to causally influence W and vice versa. Nevertheless, if they are positioned any later than that of X, the area they could influence will still be proportionately smaller than X's depending on how much later they occur.

Now, returning to Lloyd's specific example, Bunker, we can see that it is possible for Betty to causally influence the residents of Town One and the residents of Town Two, and it is also possible for the residents of Town One to influence the future residents of Town Two, but it is not possible for the residents of Town Two to influence either the residents of Town One or Betty. All of which is meant to underscore a banal (yet underappreciated) observation: *our experience of time means that we experience causation as unidirectional and are not able to influence events in the past.*¹⁵ Only presently existing people can act to promote our present interests, while people in the present and future can act to promote the interests of people in the future.

¹⁵ This may not actually be a feature of 'time', but it is an essential feature of our relationship with time.

However, Lloyd fails to draw attention to the natural corollary of the region of moral concern, which is the *region of moral dependency*—that region of spacetime that determines how things are for you right now. Just as individuals can only care about agents within their region of moral concern (and whose own region of concern thus overlaps with their own), only those in a person's region of *moral dependency* can care about them (and these people's own region of moral dependency must likewise overlap with their own).

Our view is that all the reasons for discounting for temporal location that we considered in section 2 can be shown to reduce to a single principle: *people with larger regions of moral concern and smaller regions of moral dependency have less reason to care about the wellbeing of those with smaller regions of moral concern and larger regions of moral dependency*. They merely describe four different ways in which this general principle holds true, namely that:

1. Earlier people, with larger regions of concern, have opportunities to use larger areas of spacetime to help others, while later people, with larger regions of dependency, can benefit from the resources of larger regions of spacetime.
2. The larger the discrepancy between the size of people's regions of moral concern and dependency, the less opportunities exist to engage in significant and/or reciprocal interactions that form the core of special relationships.
3. The larger someone's region of moral concern, the more agents will be existentially contingent on their choices, and the larger their region of moral dependency, the more they will be existentially contingent on others.
4. The greater the overlap between people's regions of moral concern (which will, in part, be a function of their relative size), the more they can share moral burdens.

In one sense, these four principles describe how spacetime is a resource that has value to us: It is something we can use; a medium through which we relate to one another; the means by which we shape and alter the universe; and the source of our shared moral burdens. This is why having access to more spacetime creates opportunities for growth and development by humans and other living things.

The unidirectionality of causation cannot be simulated in examples that do not involve time, for instance in ordinary examples of proximity and distance in space only, because it is a phenomenon that time does not share with spatial dimensions that we can move through in any direction. We can, however, partially simulate this by imagining cases involving a one directional portal. For instance

A mining disaster has left a group of miners injured and trapped across two mineshafts, A and B. Some of the miners in A and B are relatively free and able to help their colleagues, while other miners are so injured and/or trapped that they need their colleagues' help. Between these shafts there is a very small tunnel, which is at such an angle that it is only possible to pass through it from A to B. You are in shaft A and can help only one miner in either area. Who should you help?

In this example, while every miner is equally deserving of your help, only people who are already in A, like you, can help trapped miners in A while anyone in either A or B can help miners in B. Despite the equal value of all miners' lives, does this fact not imply that you

should have a greater degree of concern for the wellbeing of those miners trapped in shaft A than those in shaft B?

We believe it does for the following reason. Like many moral philosophers we share the sentiment that John Stuart Mill attributed to Jeremy Bentham, that “everybody to count for one, nobody for more than one” (Mill 1861, 257).¹⁶ However, to us it seems that the most important thing is that everybody is counted for one in the sense of being in receipt of an equal amount of care and moral concern, not that we ourselves must always have exactly the same degree of concern for the wellbeing of everyone everywhere across all time (cf. Parfit 1978). In the above toy case, only you and other miners in shaft A can provide care to the miners in shaft A, but the trapped miners in shaft B could receive care from people located in either one of these two mine shafts. If everyone were to treat the claims of the miners in both shafts as equally strong, then this would mean the miners in shaft B were each in receipt of more care than the miners in shaft A. The correct response to this inequality, we believe, is for the miners in shaft A to care more about the wellbeing of their trapped colleagues in shaft A than that of the miners in shaft B (note, that if these miners could move freely from one shaft to another, then this would no longer hold).

However, it is not only the mere fact of the unidirectionality of care that gives us grounds to discount the claims made on us by people in the future but also what it implies for both the number of possible recipients of care within each agent’s region of moral concern and the distribution of burdens for this care across both individuals and their access to resources.

3.b. How the numbers count

As will be clearly observed from the above, the unidirectionality of care has implications about not only who we can care for but also how many individuals we can care for. There are simply more people who can attend to the wellbeing of future people, since they can rely on aid from people in both the present and future, than can care for present people. This means that if we showed equal concern to everyone we can affect, people in the future would be in receipt of a greater total amount of care, and the benefits that might flow from this, than present people. Achieving an equitable distribution of care requires people in the present to focus their moral concern more on the present people, who only they can help. Consider the following twist on the miners case:

In one collapsed mineshaft, C, two miners can save a third, and so long as one of them takes that opportunity to help, this third miner will be saved. In another collapsed shaft, D, 200 miners can save up to 100 miners, with each miner’s help being sufficient to save another miner (but each can only save one). Again, these two areas are connected via a shaft that only allows passage from C to D.

While miners in both cases seem to have some obligation to save a fellow miner it seems to us that this duty is stronger on the first two miners in shaft C than it is on any of the miners in shaft D, and furthermore that this difference is now stronger because there are so many miners in shaft D and so few in C.

¹⁶ There is some doubt about whether Bentham ever actually said this as the only statement of Bentham’s that matches it with any degree of closeness was in a book that was largely prepared by a young J. S. Mill under his direction and supervision. See (Guidi 2008).

A weak formulation of the intuition we feel in response to this case tells us that it is simply about the numbers. While in both shafts we would have an opportunity to save a life, in shaft C, if we chose not to take up that opportunity, then we place the burden on just one other individual to save the needy miner that we refused to save from death. In the case of shaft D however, we place that burden on 199 people, many of whom (99 to be precise) could make the same choice as we did without anyone else's life being endangered. It therefore seems like we should set a higher bar for selecting inaction (e.g., a higher personal risk or cost) in the first case than we should in the second.

A stronger formulation of this intuition would draw on the idea that in the first case we seem to be in more of a unique position to aid the trapped miner, who is consequently relying on our action in a more direct and meaningful way, and this relationship produces a still stronger reason in us to rescue the third miner from death. This is similar to Larry Temkin's recent arguments (2022, chps 3-4) concerning Peter Singer's famous 'drowning child' case.

Either version of this intuition seems to imply that we have reason to be more concerned about the needs and wellbeing of people in the present, who only those now alive are in a position to help, on the basis that, even if we do not help people in the future, there are people in the future who can still do so (even if we cannot be certain they will).¹⁷ The second intuition however seems to yield a stronger result in implying that we should give additional weight to the present over the future.

3.c. Resisting the tyranny of temporal efficiency

To appreciate the full force of the significance of the unidirectionality of care however, we need to go beyond this; we must consider more than simply the potential beneficiaries of agents' choices and take account of the burdens they place upon the choosers. This is because, without applying a pure time preference, the burden of caring for the future will fall disproportionately on those in the present, even though there is a theoretically larger group of people who can share it. The reason for this is that we presently have more spacetime available, and thus more resources and potential to draw upon, when we take actions that could benefit the future. It will therefore often turn out to be more *efficient* for us to help these people than it would be for them to help themselves, as we can make investments in long term projects on their behalf. On the one hand, this seems like a valuable opportunity that we can and should use to improve the future. On the other hand, the mere fact of efficiency should not be enough to require people in the present to shoulder more than their fair share of the burdens; especially as people who live earlier in history have smaller areas of moral dependency and larger areas of moral concern to deal with.

Consider a third case:

A miner is climbing out of a collapsed mineshaft. On her way up she comes across many other miners who are trapped, and the fact that she is passing by already puts her in a good position to help them. There are people on the surface who could

¹⁷ What if we consult a crystal ball and become very confident that future people who are in a position to help less fortunate future people *will not* do so? Would we still have a stronger moral reason to aid present people? This is an important challenge to the view that we're developing here. Unfortunately, we don't have the space to develop a response. We hope to revisit it in future work. Thanks to Orri Stefánsson and Theron Pummer for raising it.

come down to aid these miners, but doing so will be harder for them than for this ascending miner.

It seems wrong to say that merely because it is easier for the one miner to save those she passes, she should have to do so unaided. It is reasonable for her to expect those on the surface to come down and aid efforts to free the trapped miners, even if that is costlier overall. Similarly, in evaluating whether we should take advantage of the efficiency gains of having access to a greater area of spacetime, it seems unfair to require present people to help the future merely because this is value maximizing.

Instead, we need to ask if this represents an especially good use of the spacetime available to them, relative to the other things they could do with it, such as helping people in the present.

One way of evaluating this is by comparing the benefits they might produce over time to the rate of interest, or returns on investment for capital. This is not because we are evaluating these benefits as a potential investment but rather because interest rates represent the average market preference for resources now relative to the future, and thus constitute a kind of collective judgement (even if merely a prudential judgement, rather than a moral one) about the value added by using these resources over a larger area of spacetime. This in turn directly reflects people's assessment of how valuable spacetime is as an asset and thus provides an overall evaluation of what they themselves see as good value for the use of spacetime.

When expected future benefits exceed present costs by more than this long-term interest rate, this indicates that this is indeed an especially good use of resources, and it is not unfair to expect present people to bare these costs. However, where it provides benefits that, while greater than the costs being paid by present people, are greater by a factor lower than the interest rate, this course of action seems to exploit the present generation for their access to spacetime, and they should not be expected to bare this cost.

This line of argument parallels what is perhaps the most widely cited argument in favour of the pure time preference for wellbeing, which was developed by Kenneth Arrow:

[Imagine] that an investment opportunity occurs, available only to the first generation. For each unit sacrificed by them, a perpetual stream of α per unit time is generated. If there were no time preference [i.e., no pure discounting], what would the optimal solution be? Each unit sacrificed would yield a finite utility loss to the first generation, but to compensate, there would be a gain, however small, to each of an infinity of generations. Thus, any sacrifice by the first generation is good. Strictly speaking, we cannot say that the first generation should sacrifice everything, if marginal utility approaches infinity as consumption approaches zero. But, we can say that given any investment, short of the entire income, a still greater investment would be preferred. I find this to be an incredible and unacceptable strain on the present generation (1999, 14).¹⁸

¹⁸ Some argue this neglects the possibility that generations might borrow for their own benefit, expecting future people to pay the bill (Parfit 1984, 484-85; Rendall 2011). In such cases, Arrow's logic might seem to support a negative discount rate, valuing future wellbeing more than present. However, Arrow is not objecting to the practise of intergenerational investment, but the moral obligations that appear to arise from the efficiency of investing to benefit later people, and there is no analogous argument in favour of borrowing.

One way to understand the injustice of such efficiency-based obligations is by way of what Larry Temkin calls the ‘disburse additional burdens view’ (2022, 239). This is the claim that it is unfair for any individual or group of individuals to suffer great costs merely for others to each obtain a smaller benefit, even if the group who stands to benefit is much larger than that which stands to lose out.¹⁹

Notice how this argument interacts with a concern for equality. It can be easy to assume that a concern for a fair distribution of benefits and burdens and a concern about equality should lead to the same conclusions. But they need not, as Parfit, Temkin, and others have pointed out. In many cases these principles seem to come apart, and if we want to produce a more equal distribution, this will entail disproportionately large sacrifices on the part of a smaller group of well-off people with smaller benefits being spread over a larger number of those who are worse off. However, regarding time there is no such general conflict. The same facts about our relationship to spacetime that tend to make it more efficient for early generations to act for the benefit of later ones also make it more likely that later generations will be better off.²⁰ This coincidence highlights once more how divergent moral intuitions can point to the moral significance of humanity’s relationship with spacetime.

4 | Existential Risk and Astronomical Waste

It can sometimes seem like an article of faith among those who are concerned about the future survival of our species that we must reject the practise of pure time discounting. A key source for this widely held belief is an appendix to Parfit’s influential book *Reasons and Persons*, in which he briefly surveys and rejects six arguments that have been used to defend the practise (two of which we considered in sections 2a and 2c). Another was a prominent debate between Stern and Nordhaus following the publication of Stern’s *The Economics of Climate Change*, in which both parties agreed that the lower the discount rate the stronger the argument would be in favour of climate change mitigation.

The debate between Nordhaus and Stern is instructive. The basis for the argument is that both these economists conceptualize climate change in the following way:

Another objection is that Arrow’s argument wrongly implies that, if we must choose between two, otherwise inevitable, catastrophes, one sooner but more minor and the other later and more severe, we should avert the earlier catastrophe. However, this theoretical possibility is unlikely to the point of irrelevance. Far more likely is that only we can avert the earlier catastrophe, while future generations will have some chance to avert the later one even if we do not. Suggesting that we must avert the later catastrophe in this case, because this is the more efficient use of our resources, not only dooms people to suffer the earlier catastrophe but lets future generations off the hook for averting the later catastrophe.

¹⁹ For discussion, see (Temkin 2014). Other principles implying similar conclusions include Parfit’s Wide Dual Person Affecting Principle (Parfit 2017, 154) and the ‘Competing Claims’ view (Otsuka and Voorhoeve 2009; Voorhoeve and Fleurbaey 2012; Otsuka 2018).

²⁰ On the other hand, if humanity is expecting an oncoming catastrophe, then future generations might well end up worse off than we are, giving us an egalitarian reason to have more concern for their wellbeing. It is also likely that, in such cases, present sacrifices will be an especially effective means of securing future wellbeing, as we work to either prevent or mitigate the severity of this catastrophe. Depending on the circumstances these two factors might effectively cancel each other out; they might give us reason to adopt a higher pure time preference; or they might give us reason to adopt a negative discount rate, valuing future wellbeing more than present (albeit for the same egalitarian reasons we standardly value it less).

In the future, we will experience continuous economic growth for the rest of time; however, if we act to mitigate climate change, then the trajectory of this growth will be less in the short term and more in the long term, while if we don't mitigate climate change, then the economy will grow more in the short term but less in the long term.

Of course, if one believed that this was the choice we faced in responding to climate change, then one would only see mitigation as a good thing if one was relatively indifferent between wellbeing in the short and long term. However, many people would not see this as the climate policy choice we face (Rendall 2019; Shue 2021), and it certainly does not describe the reality of existential risk. Our futures could be much worse than our present, and it may be that climate change, or other phenomena, will mean that humanity has no future at all. Indeed, the idea that we should make sacrifices in the name of economic growth (described here as climate mitigation) may be precisely what is causing climate change.²¹

Instead, this trade-off presented in the debate between Nordhaus and Stern is similar to another argument that has been made around existential risk, Nick Bostrom's (2003) argument against astronomical waste. On this argument, we should do whatever we can to move towards a state of 'technological maturity' (in which posthumans colonize the stars) as quickly as possible, because the amount of wellbeing that might exist in such a state is so much greater than anything we could experience in our present condition, that bringing forward the point at which we transition would be worth almost any sacrifice on the part of presently existing people.

We do not wish to take a view on whether these two arguments are either morally compelling or accurately describe the possible future trajectories facing us. Instead, we bracket all considerations about what kind of future would be best for humanity, and ask whether it is necessary to oppose the pure time preference for wellbeing, if one is concerned about decisions that may affect whether humanity has any future at all?

In general, discounting the value of future wellbeing will significantly reduce our moral obligations to make the future only marginally better or worse than it would have been, even if relatively small differences in the amount of future wellbeing could still be huge (in absolute terms) compared with some relatively much larger sacrifice from presently existing people (for instance a loss of 0.01% of future wellbeing in 100,000 years may come to the same amount as a loss of 50% of wellbeing today). However, it doesn't mean that we don't need to care about changes in the quantity of future wellbeing that are large in both absolute terms and relative to the expected quantity of future wellbeing, such as whether any people will exist in the future at all or whether they experience a very good future or a future of mere subsistence.

Yes, a pure time preference will reduce the strength of our obligations to mitigate existential risk relative to doing nothing to benefit future people, to some degree. However, this reduction will be relatively small and unimportant. On the other hand, introducing a pure time preference for wellbeing will mean that our obligations to ensure that humanity has a future become far more important than our obligations to maximize the sum of future wellbeing. In this sense, these arguments may leave us with even more reason to prevent

²¹ This point was made in the subsequent Dasgupta Review of the Economics of Biodiversity (Dasgupta 2021).

the real-world existential catastrophes facing humanity, many of which stem from our headlong rush towards technological maturity, than we had before.

Acknowledgments

For their feedback on this paper and the ideas contained within it we are grateful to Partha Dasgupta, Orri Stefánsson, Matthew Rendall, Harry Lloyd, J. Paul Kelleher, Theron Pummer, Niels Brøgger, an audience at the workshop Climate Ethics and Climate Economics: Discounting, what have we learned? at the University of Nottingham, and three anonymous reviewers. This paper was made possible via support from the Templeton World Charity Foundation; the views expressed are those of the authors and do not necessarily reflect the views of the foundation.

Bibliography

Arrow, Kenneth J. (1999). Discounting, Morality, and Gaming. In P. R. Portney and J. P. Weyant (eds.), *Discounting and Intergenerational Equity* (pp. 13-22). Washington, DC: Resources for the Future.

Beard SJ, and Kaczmarek Patrick (forthcoming) Evaluating Extreme Technological Risk, In C. R. Rhodes (ed), *Managing Extreme Technological Risk*. World Academic Publishing.

Bleske-Rechek, A., Nelson, L. A., Baker, J. P., Remiker, M. W., & Brandt, S. J. (2010). Evolution and the Trolley Problem: People Save Five Over One Unless the One is Young, Genetically Related, Or a Romantic Partner, *Journal of Social, Evolutionary, and Cultural Psychology* 4(3): 115-127.

Bostrom, Nick (2003). Astronomical Waste: The Opportunity Cost of Delayed Technological Development, *Utilitas* 15(3): 308-314.

Caney, Simon (2014). Climate Change, Intergenerational Equity and the Social Discount Rate, *Politics, Philosophy & Economics* 13(4): 320-342.

Cotton-Barratt, Owen, Farquhar, Sebastian, Halstead, John, Schubert, Stefan, and Snyder-Beattie, Andrew (2016). Global Catastrophic Risks 2016, *Global Challenges Foundation*, <URL: <http://globalprioritiesproject.org/2016/04/global-catastrophic-risks-2016/>>.

Dasgupta, Partha (2019a). Ramsey and Intergenerational Welfare Economics. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Summer 2020, <URL: <https://plato.stanford.edu/archives/sum2020/entries/ramsey-economics/>>.

— (2019b). *Time and the Generations: Population Ethics for a Diminishing Planet*. New York: Columbia University Press.

— (2021). *The Economics of Biodiversity: the Dasgupta review*. Cambridge: Cambridge University Press

- Heath, Joseph (2021). *Philosophical Foundations of Climate Change Policy*. Oxford: Oxford University Press.
- Greaves, Hilary (2017). Discounting for Public Policy: A Survey, *Economics & Philosophy* 33(3): 391-439.
- Guidi, Marco (2008). Everybody to Count for One, Nobody for More than One: The Principal of Equal Consideration of Interests from Bentham to Pigou, *Revue D'études Benthamiennes* 4: 40-69.
- Hutchinson, Michelle (2016). People Aren't Replaceable: Why It's Better to Extend Lives than to Create New Ones. In E. Gamlund and C. T. Solberg (eds.), *Saving People from the Harm of Death* (pp. 203-214). New York: Oxford University Press.
- Kaczmarek, Patrick and Beard, SJ (2020). Human Extinction and Our Obligations to the Past, *Utilitas* 32(2): 199-208.
- Kamm, F. M. (2004). Failures of Just War Theory: Terror, Harm, and Justice, *Ethics* 114(4): 650–692.
- Kelleher, J. Paul (2017). Pure Time Preference in Intertemporal Welfare Economics, *Economics & Philosophy* 33(3): 441-473.
- Lloyd, Harry (ms). Time Discounting, Consistency and Special Obligations: A Defense of Robust Temporalism, *Global Priorities Institute (Working Paper No. 11-2021)*, <URL = <https://globalprioritiesinstitute.org/time-discounting-consistency-and-special-obligations-a-defence-of-robust-temporalism-harry-r-lloyd-yale-university/>>.
- Mill, John S. (1861). Utilitarianism. In J. M. Robson (ed.), *The Collected Works of John Stuart Mill, Volume X, Essays on Ethics, Religion and Society* (pp. 203-259). Toronto: University of Toronto Press and Routledge & Kegan Paul.
- Mogensen, Andreas (2022). The Only Ethical Argument for Positive δ ? Partiality and Pure Time Preference, *Philosophical Studies* 179(9): 2731-2750.
- Narveson, Jan (1973). Moral Problems of Population, *The Monist* 57(1): 62-86.
- Ng, Yew-Kwang (2016). The Importance of Global Extinction in Climate Change Policy, *Global Policy* 7(3): 315-322.
- Nordhaus, William (2007). Critical Assumptions in the Stern Review on Climate Change, *Science* 317(5835): 201-202.
- Otsuka, Michael (2018). How It Makes a Moral Difference That One Is Worse Off Than One Could Have Been, *Politics, Philosophy & Economics* 17(2): 192-215.

Otsuka, Michael and Voorhoeve, Alex (2009). Why it Matters That Some Are Worse Off Than Others: An Argument Against the Priority View, *Philosophy & Public Affairs* 37(2): 171-199.

Parfit, Derek (1978). Innumerate Ethics, *Philosophy & Public Affairs* 7(4): 285-301.

— (1984). *Reasons and Persons*. Oxford: Oxford University Press.

— (2016). Can We Avoid the Repugnant Conclusion?, *Theoria* 82(2): 110-127.

— (2017). Future People, the Non-Identity Problem, and Person-Affecting Principles, *Philosophy & Public Affairs* 45(2): 118-157.

Pigou, Arthur C. (1932). *The Economics of Welfare (4th Edition)*. London: Macmillan & Co.

Ramsey, Frank P. (1927). A Contribution to the Theory of Taxation, *The Economic Journal* 37(145): 47-61.

— (1928). A Mathematical Theory of Saving, *The Economic Journal* 38(152): 543-559.

— (1931). Epilogue. In R.B. Braithwaite (ed.), *The Foundations of Mathematics and Other Logical Essays*. London: Routledge and Kegan Paul.

Rendall, Matthew (2011). Climate Change and the Threat of Disaster: The Moral Case for Taking Out Insurance at our Grandchildren's Expense, *Political Studies* 59(4): 884-899.

— (2019). Discounting, Climate Change, and the Ecological Fallacy, *Ethics* 129(3): 441-463.

Shue, Henry (2021). *The Pivotal Generation: Why We Have a Moral Responsibility to Slow Climate Change Right Now*. Princeton: Princeton University Press.

Stern, Nicholas (2006). *Stern Review: The Economics of Climate Change*. Oxford: Oxford University Press

Stern, Nicholas and Taylor, Chris (2007). Climate Change: Risk, Ethics, and the Stern Review, *Science* 317(5835): 203-204.

Temkin, Larry S. (2014). *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. Oxford: Oxford University Press.

— (2022). *Being Good in a World of Need*. Oxford: Oxford University Press.

Voorhoeve, Alex and Fleurbaey, Marc (2012). Egalitarianism and the Separateness of Persons, *Utilitas* 24(3): 381-398.