Philos Stud (2022) 179:2731–2750
https://doi.org/10.1007/s11098-022-01792-8

The only ethical argument for positive δ? Partiality and pure time preference

Andreas L. Mogensen[1]

Accepted: 11 January 2022 / Published online: 4 February 2022
© The Author(s) 2022

Abstract I consider the plausibility of *discounting for kinship*, the view that a positive rate of pure intergenerational time preference is justifiable in terms of agent-relative moral reasons relating to partiality between generations. I respond to Parfit's objections to discounting for kinship, but then highlight a number of apparent limitations of this approach. I show that these limitations largely fall away when we reflect on social discounting in the context of decisions that concern the global community as a whole, such as those related to global climate change.

Keywords Discounting · Climate change · Partiality · Agent-relativity

1 Introduction

What costs should the current generation be willing to bear for the sake of those that follow? Suppose we understand costs and benefits as changes in *generalized consumption*, taking this to include not only goods and services exchanged in the economy, but also non-market goods such as leisure and health. Under what circumstances is it desirable to reduce consumption in the current generation so as to increase consumption in future?

The aim of this paper is to help answer this question by considering whether a *positive rate of pure intergenerational time preference* is justifiable. Roughly speaking, a positive rate of pure intergenerational time preference means that we care more about the welfare of people insofar as they are nearer to us in time. Broome (1994, 128) writes: "To many philosophers this seems a reprehensible practice." In this paper, I argue that a positive rate of pure intergenerational time

✉ Andreas L. Mogensen
andreas.mogensen@philosophy.ox.ac.uk

[1] Global Priorities Institute, Trajan House, Mill Street, Oxford OX2 0DJ, UK

preference is nonetheless justifiable as a means of representing agent-relative reasons relating to partiality between generations.

The next section explains the idea of a social discount rate on investment projects funded by reductions in current consumption. Section 3 responds to two objections raised by Parfit (1984) against appealing to agent-relative reasons to justify the use of a social discount rate. In Sect. 4, I outline three further observations about the apparent limitations of agent-relative discounting. Each turns on the fact that any bonds of partiality that diminish with time apparently link us to only a proper subset of all people, whether now or in future. In Sect. 5, I argue that these limitations largely drop away when we reflect on cases like global climate change, where the decisions at issue concern the global community as a whole. Section 6 concludes.

2 Pure time preference

Under what circumstances is it desirable to reduce consumption in the current generation to increase consumption in future? The *social discount rate* is the standard tool used by economists to answer this question. This section explains the nature of the discount rate, outlines the Ramsey formula, and describes existing debates about the rate of pure time preference in ethics and economics.

2.1 The social discount rate

The social discount rate is the minimum (riskless) rate of return that must be earned by an investment project in order for its implementation to be socially desirable (Gollier, 2013). Ramsey (1928) specifies a well-known formula for determining the discount rate, so understood: $r = \delta + \eta g$. Here, r is the discount rate, and g is the *growth rate of consumption*. The remaining parameters are so-called 'taste parameters.' Thus, η is the *elasticity of the marginal utility of consumption*, typically understood as a measure of our aversion to intertemporal consumption inequality. Lastly, the δ parameter specifies the *rate of pure time preference*. This is the proportional rate of decline in the *utility discount factor*, which is the weight we put on utility derived from consumption occurring at a given point in time. The utility discount factor is typically normalized to 1 in the current time period. If its value declines as a function of time, we have a positive rate of pure time preference: we care more about utility derived from consumption if it occurs sooner, rather than later.

The choice of discount rate matters greatly for how we think of our obligations to posterity. The disagreement between Nordhaus (2007, 2008) and Stern (2006, 2008) concerning the extent of emissions abatement required in the near future is traceable principally to conflicting views about discounting. This disagreement is driven largely by differences of opinion concerning the values of the two 'taste parameters.' For example, Nordhaus adopts a value for δ of 1.5%, whereas Stern rejects a positive rate of pure time preference as ethically indefensible, paring δ down to a pure measure of the exogenous per year risk of human extinction, which he sets at 0.1%.

Nordhaus (2007, 692) attempts to paint Stern's commitment to intergenerational impartiality as controversial and idiosyncratic, alleging that it "stems from the British utilitarian tradition with all the controversies and baggage that accompany that philosophical stance." However, Stern's view is widely shared among moral philosophers, including those hostile to utilitarianism. For example, Rawls (1999, 253) tells us, in no uncertain terms, that "there are no grounds for discounting future well-being on the basis of pure time preference".

2.2 Pure time preference and agent-relative reasons

The basic moral argument against a positive rate of pure time preference is stated clearly in a recent paper by Caney (2014, 323–324)[1]:

> A person's place in time is not, in itself, the right kind of feature of a person to affect his/her entitlements. For example, it does not make someone more or less deserving or meritorious. Similarly, it does not, in itself, make anyone's needs more or less pressing. … It is not the right kind of property to confer on people extra or reduced moral status.

What should we make of this line of argument?

I want us to focus on the following observation. Generally speaking, in order to be justified in caring more about some people than others, we need not regard the latter as having 'reduced moral status.' For example, I care more about my wife than any of you reading this paper, but I don't regard you as having a lower moral status than she does. I do not think there is a general requirement on moral agents to value my wife's well-being more highly than yours, for example. Thus, in general, in order to be justified in caring more about the utility of people who are nearer to us in time and less about those who are more distant, we need not regard location in time as conferring on people extra or reduced moral status.

We can re-state this point, appealing to the familiar distinction between *agent-neutral* and *agent-relative* moral reasons (Nagel, 1970; Parfit, 1984). Caney's argument appears to concern only the question of whether there are agent-neutral moral reasons for assigning lower weight to the utility of person 1 than person 2, such as that person 2 is more deserving than person 1, or is otherwise imbued with higher moral status. But some of the most important moral reasons that we have for caring more about one person than another are agent-relative reasons having to do with the particular relationships in which we stand to some people and not to others.

It may be suggested that the neglect of agent-relative reasons is justified in this context because the social welfare function from which the Ramsey formula is derived is agent-neutral, representing an impartial welfarist valuation of consumption streams, as opposed to a valuation that takes up the perspective of any particular individual (compare Intergovernmental Panel on Climate Change 2015). It is indeed common to describe this social welfare function as 'utilitarian'. However, as Kelleher (2017) argues, this is an artefact of economists' decision to retain the

[1] Caney also rejects utilitarianism: see Caney (2009, 168).

mathematical structure of Ramsey's theory. The utilitarian moral theory that Ramsey presupposed is not necessarily retained in addition, and many economists treat the social welfare function in ways that make allowance for the existence of agent-relative reasons. For example, Dasgupta (2008) argues that our choice of the value of η should be influenced by whether or not *we* are responsible for any inequalities that fall within the scope of evaluation.

In reminding ourselves of the moral significance of partiality (Keller, 2013; Kolodny, 2010; MacIntyre, 1984; Scheffler, 1997, 1999, 2004; Williams, 1982; Wolf, 1992), can we construct a plausible rationale for a positive rate of pure intergenerational time preference? This idea has been raised by a number of economists (Arrow, 1996; Beckerman and Hepburn, 2007; Schelling, 1995; Stern, 2008). Stern (2008, 15) describes it as the "only … ethical argument for positive δ … that has some traction."[2] However, aside from the objections levelled against it by Parfit (1984, 485–486) (recapitulated in Cowen and Parfit, (1992) and in Broome (1992)), this idea has received relatively little critical discussion among philosophers.[3]

3 Parfit on discounting and partiality

The neglect of agent-relativity as a justification for pure time discounting among philosophers might, of course, be credited to the fact that Parfit refuted the idea. In this section, I set out and respond to Parfit's objections. I begin by characterizing the particular variant of this idea that Parfit discusses. I then consider the two objections raised in Parfit's discussion, responding to each in turn.

3.1 Discounting for kinship

The particular variant of the idea that a positive rate of pure intergenerational time preference can be justified in terms of agent-relative reasons discussed by Parfit is a view I will call *discounting for kinship*. The basic idea is as follows. The people who are born into the next generation are our children. (Or, if we are now childless and expect to remain so, then they may be, say, our nephews and nieces.) By the lights of common-sense morality, we are each permitted and/or required to be

[2] This is not to say that Stern accepts the argument. He suggests that partiality toward one's nearest and dearest "is often explained in terms of functionality for survival of groups. However, this type of reasoning from evolutionary biology does not have much relevance when we are thinking about the survival of the planet as a whole." (15).

[3] The recent paper by Kelleher (2017) also offers a partial defense of agent-relative discounting against some objections to pure time discounting raised by Broome (2004, 2012), but his primary focus is showing that economist's valuations of consumption streams do not presuppose an agent-neutral, axiological evaluation of outcomes. Kelleher concludes that "we need further substantive discussion of intergenerational ethics to decide whether pure [agent-relative] discounting should be a part of the current generation's response to the situation it faces." (469) Another relevant antecedent of the general approach to be explored in this paper is the communitarian approach adopted by De-Shalit (1995), which I note in Sect. 3.1.

strongly partial to the interests of our children (and also, presumably, to our nephews and nieces). Plausibly, we are also permitted and/or required to be partial to our grandchildren (and also to our grandnephews and grandnieces),[4] but to a lesser degree. For each succeeding generation, the degree of permissible and/or obligatory partiality declines as the closeness of the bonds linking present and future people decline. Therefore, we can permissibly weight the welfare of each succeeding generation less than that of the generation preceding it.[5]

This view should not be understood as committing us to the idea that a merely biological relationship justifies partiality among kin. I will treat 'kinship' as referring to whatever relationship among those individuals whom we recognize as family members is associated with reasons for such individuals to be partial to one another. Some may think that biology plays an important role here (e.g., Kolodny, 2010; McMahan, 2002, 375–377; Velleman, 2005, 2008). Others will deny this (e.g., Boonin, 2003, 227–234). So far as I can see, discounting for kinship does not require us to take a stand in this debate.[6]

Obviously, it is not the mere passage of time that justifies the use of a declining utility discount factor according to this view. What matters ultimately is not distance in time, but in the family tree. However, in practice it will generally be true that the more distant some one of our descendants is from us in time, the more distantly related are we. As a simplifying idealization or heuristic, it therefore seems warranted to represent discounting for kinship in terms of a declining utility discount factor when constructing our social welfare function.

Some may object that this means we are not discussing *genuine* pure time discounting, since no fundamental moral importance is attached to the passage of time. That may be so. It may well be the case that the philosophical consensus against the use of positive δ is merely a consensus that no basic ethical significance should be attached to the movement of time. But if so, then the philosophical consensus arguably has a less obvious bearing on the practical questions that divide economists. Economic models used in the analysis of optimal climate policy include

[4] For ease of exposition, I henceforth omit these parenthetical remarks about oblique kinship relations and speak as if only vertical kinship relations were significant in linking us to future generations.

[5] It may already have occurred to some readers that a significant problem for this view arises from the fact that only a very small fraction of each subsequent generation is closely related to any one person, raising questions about how this approach could in principle justify discounting the welfare of subsequent generations as a whole. This issue is discussed at length in Sects. 4 and 5 of the paper.

[6] Similarly, I believe we need take no stand on the ultimate philosophical justification for partiality. With respect to the influential tripartite taxonomy proposed by Keller (2013), I believe the position I develop to be consistent with the Projects View, the Relationship View, and the Individuals View, as well as syncretic views that combine elements of each (see Lord, 2016). Admittedly, it is sometimes unclear whether a particular philosophical justification of partiality will support the claims that I present here. Consider, in particular, a rule-consequentialist justification for partiality (Parfit, 2011, 385; Rendall, 2019, 448–449). The scope of permissible partiality is harder to pin down on this view, because it depends so heavily on empirical contingencies. The difficulties are especially stark in contexts where long-run impacts are salient, such as climate change. Thus, the views I develop are arguably more naturally located within a moral perspective that assigns a more foundational and less contingent role to partiality as a source of moral reasons, although they are not inconsistent in principle with a rule-consequentialist approach.

a wide range of idealizing simplifications. They are models, after all. A declining utility discount factor may be understood as yet another idealizing simplification, which could be taken to represent discounting for kinship.

Here is one final point before we move on. While discounting for kinship is the focus of discussion in this paper, structurally similar approaches can arguably be constructed by highlighting other relationships associated with special obligations, whose significance decays gradually over time.[7] For example, in his communitarian theory of intergenerational justice, De-Shalit (1995, 58) argues that special obligations that arise from membership in identity-constitutive communities fade over time: "as our community progresses through time, it tends to reach a lesser and lesser degree of cohesion. Since obligations derive from relationships and a sense of community, and since these are diminished the obligations are diminished also."[8] I will nonetheless focus on discounting for kinship. In part, it is just useful to have a concrete focus, rather than to talk about partiality in the abstract. Moreover, the existence of very powerful special obligations associated with ties of kinship is widely accepted. De-Shalit's approach, by contrast, is developed against the background of a controversial communitarian political philosophy.

3.2 Discounting harms

What, then, are Parfit's objections to using discounting for kinship to justify "employ[ing] a standard Discount Rate." (485)?

On the one hand, Parfit insists that discounting for kinship cannot justify the use of a discount rate "to the infliction of grave harms." (486) His view is that when it comes to the imposition of such harms, "special relations make no moral difference." (ibid.) Therefore, even if we insist that diminishing bonds of kinship can justify some kind of diminishing concern for future persons, it cannot justify the application of a declining utility discount factor "to all kinds of effect." (ibid.)

The view expressed by Parfit here bears a notable resemblance to what Caney (2009) calls the *Scope Restricted View*. The Scope Restricted View says that violations of human rights are not to be discounted for time, although we may discount other costs and benefits. Parfit's claim that special relationships are morally irrelevant when it comes to the infliction of grave harms may be thought to imply that discounting for kinship entails the Scope Restricted View.

Even if we grant that the only reasons against inflicting grave harms are agent-neutral, the significance of this observation may be limited in light of the Non-Identity Problem (Parfit 1984, 351–379). When it comes to significant and far-reaching investment projects, any choice we make will have significant and far-reaching effects on the identities of future people. A few centuries hence, there may be no person impacted by our decision who would have existed had we chosen differently. If a counterfactual-comparative account of harm is correct and a person

[7] I'm grateful to an anonymous referee and to Adam Bales for pressing me on this point.

[8] I'm grateful to Harry Lloyd for bringing my attention to De-Shalit's views on this point. Note that De-Shalit's remarks on discounting per se are highly critical (De-Shalit 1995, 84–85).

cannot be harmed (in the morally relevant sense) by some action unless they are made worse off than they would have been otherwise (Parfit 1984, 69; Jackson 1997; Norcross 2005; Boonin 2014), that means these people cannot be harmed (in that sense) by our choice, since they are not worse off than they would have been had we chosen differently. Therefore, there may be no objection along the lines described by Parfit to discounting the pain and suffering experienced by these people as a result of climate change.

To be clear, the suggestion here is not the absurd idea that when our actions today change the identities of future people, the impact of our actions on the welfare levels of future people is morally indifferent. Rather, it is the far more modest point (which Parfit would surely accept) that anyone who puts special emphasis on *inflicting harm* as a morally significant category of analysis in thinking about ethics and future people will face serious challenges in analysing cases that invoke the Non-Identity Problem.[9]

In any case, there are additional reasons to resist Parfit's objection. In particular, I think it is clear that special relations *do* make an important difference to the morality of inflicting grave harms. In Dante's *Inferno*, there is literally a special place in Hell for people who kill their parents, siblings, or children. Although it may be permissible to turn the trolley from the five and onto the one in the *Switch* case (Foot 1967; Thomson 1985), intuitively, your reasons against doing so would be much stronger if the one on the side-track were your own child, and we may well think it would be *im*permissible to turn the trolley in this case, as Kamm (2004, 674) does.[10] Drawing an analogy with hypothetical life-or-death choices of this sort, Kamm (2004) and Hurka (2005) both argue that in war contexts national partiality justifies discounting collateral damage to enemy non-combatants relative to collateral damage one might inflict on one's compatriots. Whether this analogy is acceptably

[9] Of course, there are many philosophers who believe these challenges can be met by rejecting or modifying the counterfactual-comparative account of harm (see, *inter alia*, Hanser 2008, 2009; Harman 2004, 2009; Shiffrin 1999; Woollard 2012). In my view, there are good reasons to doubt the plausibility of these alternative theories. Consider, for example, Harman's view, on which, roughly, to harm someone is to cause them to be in a bad state. Imagine that Paul is being tortured and experiences a pain of intensity 10. The only thing I can do to lessen Paul's pain is to pull a switch so that he will experience a pain of intensity 1 instead. On Harman's view, it seems, absurdly, that I harm Paul by pulling the switch, since my act causes him to be in a bad state: i.e., experiencing pain of intensity 1. Of course, this objection is far from knock-down. Unfortunately, there is not space within this paper to further explore this controversy. See Boonin (2014, 52–102) for further discussion.

[10] Consider also this variant of the *Switch* case, due to Lackey (1984, 165): "The trolley is careening down the tracks. If the switchman does not throw the switch, the trolley will go forward and run over the switchman's son. If the switchman throws the switch, the trolley will go onto a spur and kill a stranger. Many people feel that it would be permissible for the switchman to throw the switch". I'm grateful to an anonymous referee for bringing my attention to this example. As the same referee notes, intuitions about permissible harm and partiality can vary on the basis of details whose moral significance may be unobvious or even suspect. Thus, in contrast to our intuitions about the *Switch*-style case, we think it morally impermissible for a mother to strip a lifejacket from one child in order to save her own from drowning instead (Lackey 1984, 161). *Prima facie*, this is just the general problem of giving a coherent and intuitively satisfying account of deontological constraints on harming. Thus, Lackey suggests that our judgments about the cases noted here can be rationalised by something like Thomson's 'distributive exemption' (Thomson 1985).

drawn is controversial, since we may doubt that reasons of partiality among compatriots have the same kind of force in war contexts as reasons of partiality among kin in *Switch*-style cases (see Lefkowitz 2009). Even so, these *Switch*-style cases show that the view that special relations make no difference at all to the morality of imposing grave harms lacks intuitive force.[11]

3.3 Constant δ

The other objection noted by Parfit in arguing that discounting for kinship cannot justify the use of 'a standard Discount Rate' is that if we discount for kinship, then δ must eventually drop to or asymptotically approach zero. We must keep in mind that what is being discounted for time, on the present view, is the greater weight that attaches to the interests of certain people in virtue of their kinship to us. Therefore, the practice of discounting for kinship should not lead us to count the welfare of some descendant of ours for less than that of a complete stranger living now. As Parfit says: "We ought to give *some* weight to the effects of our acts on mere strangers. We ought not to give *less* weight to effects on our own descendants." (486)[12] However, if we discount the utility of our descendants using a constant positive rate of pure time preference, we end up violating this constraint.

This suggests a second key factor that Parfit thinks of as constituting 'a standard Discount Rate': namely, the use of constant δ. This assumption is indeed standard in cost–benefit analysis. It is typically assumed that social preferences should be *dynamically consistent*: roughly speaking, this means that a plan must not be socially preferred at one point in time but socially dispreferred at a later point, assuming no new information about the decision problem is acquired in the intervening period. A declining rate of pure time preference violates this constraint (Strotz 1955).

Perhaps all Parfit intends to argue is that discounting for kinship cannot justify the use of constant δ. If so, this fails to show that discounting for kinship cannot justify the use of a declining utility discount factor whose proportional rate of

[11] Someone who concedes that reasons of partiality related to kinship can make an important difference to the morality of imposing grave harms in some cases but wants to insist that they fail to do so in the context of discounting harms due to climate change would owe us some explanation for why this should be so. No such explanation is suggested in Parfit's discussion. An anonymous referee suggests that if we have a choice between imposing a significant cost on x or on y, our reasons to impose the cost on x as opposed to y are stronger if y is our kin; but if we also have the option of suffering a smaller personal cost to avoid imposing a significant cost on x or y, then our reasons not to impose a cost on y may be no stronger than our reasons not to impose an equivalent cost on x. This is an interesting suggestion, but it leaves a significant explanatory burden in place. It is in general surprising to be told that our reasons to choose a over b depend on the availability of a third option, c, and someone advancing this sort of claim typically owes us an explanation (compare Voorhoeve 2014, 78–79). Thus, in the present context, we need some explanation for why the relative strength of my reasons to impose the cost on x as opposed to y weaken due to the availability of the third option.

[12] Readers who have anticipated later developments in the paper may already have formulated a reply to Parfit's objection, appealing to the fact that if the rate of pure time preference models the rate at which we as a whole should discount the welfare of future generations as a whole, then there are no presently existing people who are strangers outside the presently existing collective, since it is all-encompassing. This reply is discussed in Sect. 5.5.

decline is itself declining.[13] However, the point about dynamical consistency noted above may seem to make the objection more general and more forceful.

So understood, Parfit's second point can be thought of as related to a criticism that Broome (2004, 2012, 2016) has raised against pure time discounting. Broome notes that we could in principle engage in *temporally neutral* pure time discounting: we could think that the welfare of twenty-first century people should be assigned greater weight than the welfare of twenty-second century people by anyone at all, no matter their position in time. In practice, Broome notes, pure time discounting is interpreted as involving a *time-relative* value function: the utility discount factor assigned to a person's welfare differs depending on what time the decision maker occupies. This suggests that earlier and later valuations of consumption streams made by one and the same decision-maker may conflict in ways that are foreseeable. According to Broome (2016, 909), "This sort of incoherence in policymaking is not consistent with rational agency."

In my view, the claim that dynamic consistency is a general desideratum on social preferences is far from compelling (Heal 1998; Henderson and Bateman 1995; Hepburn 2006). It seems especially implausible when a change in social preference is explained by changes in the composition of society, such that later generations have agent-relative reasons to weight the interests of people differently than earlier generations.In general, a charge of inconsistency seems applicable only to attitudes that belong together to a suitably unified outlook. If I believe both p and $\neg p$, you may accuse me of being inconsistent. If I believe p and you believe $\neg p$, we together cannot be charged with the rational failing of inconsistency. If we think that individual persons are suitably unified over time, then dynamic consistency may represent a rational requirement on individual choice behaviour.[14] Social collectives, being successively constituted by wholly distinct people with different loyalties and different preferences, certainly do not have the degree of unity across time needed to make a general requirement of time consistency plausible.

The claim for which I am arguing here is, I think, ultimately compatible with Broome's views. Broome (2016, 909) emphasizes that what he objects to is dynamic inconsistency at the level of the individual. The incoherence on which he focuses arises when one and the same agent successively occupies distinct times and her preferences over plans are represented by a time-relative value function, leading to dynamic inconsistency. Discounting for kinship does not entail that any individual's own preferences are to be represented by a time-relative value function. In discounting for kinship, we put more weight on the utility of our grandchildren than our great-great-grandchildren. In practice, that means we care more about

[13] On declining discount rates in the economic evaluation of public policy, see Greaves (2017, 417–418). Jamieson (2014, 128) notes that, in general, a declining pure rate of time preference "seems to conform more closely to how people think and behave than does a fixed value for *delta*." Note that strictly speaking, a declining rate of pure time preference is insufficient to avoid Parfit's objection. In addition, we require that the discount factor asymptotically approaches a positive limit. I'm grateful to an anonymous referee for pressing me to clarify this point.

[14] In fact, Parfit's own views about the (dis)unity of the person entail that prudentially rational preferences may be dynamically inconsistent (Ahmed 2020).

temporally nearer generations. However, holding fixed that the valuation is made from our perspective and from within the network of relationships within which we find ourselves, it is invariant with respect to our position in time. No matter where we might be transported by a Wellsian time machine, we would still be more closely related to our grandchildren than to our great-great-grandchildren and would assign relative weights to their utility accordingly. At the individual level, there is no scope for foreseeable preference reversals induced by the mere passage of time. This occurs only at the level of social preferences, when later generations have agent-relative reasons to weight the interests of people differently than earlier generations. While social decision-makers should obviously take this possibility into account when setting policy, they need not otherwise fear that they are choosing irrationally. For this reason, I view discounting for kinship as ultimately immune to Broome's critique of pure time discounting.

Assuming that we find this persuasive, Parfit's second objection loses its punch. It may be true that the standard practice of discounting assumes constant δ. And it may be true that adopting a declining rate of pure time preference leads to dynamic inconsistency. However, time consistency is not a plausible constraint on reasonable social preferences, for the reasons I have noted, so there is no general objection to discounting for kinship in play here.

4 Additional problems

I have defended discounting for kinship against Parfit's objections. However, there are other objections to consider. In this section, I raise three observations about the apparent limitations of discounting for kinship. Each observation turns, in its own particular way, on the fact that any bonds of partiality that diminish with time seem to link us to only a proper subset of all people, leaving us with strangers in our midst, both now and in future.

4.1 First limitation

Recall the observation that the application of a declining utility discount factor to the well-being of our descendants should never lead us to value the welfare of some descendant of ours less than that of some currently existing stranger. Plausibly, we should never end up in a position where we value the life of some member of the current generation of our kin over those of n distantly related descendants existing at time t unless we also value one of our family members living now more than n currently existing strangers. If the decline of the discount factor is intended to mimic the fact that we are permitted to assign greater weight to our near kin than to our distant descendants, we seem to be constrained in this way in setting the schedule of discount factors. Furthermore, observing this constraint calls into question what might otherwise seem like relatively modest suggestions for intergenerational pure time discounting.

To see this, suppose we are wondering how much weight should be put on the utility of our descendants living 500 years from now. We might wonder whether it

could be appropriate to weight their welfare as if we had applied a constant rate of pure intergenerational time preference of 1.5%, compounded annually. Of course, we know from our discussion in Sect. 3.4 that we should not use a constant rate of pure time preference. Nonetheless, we may wonder whether the application of the proper declining rate would lead us to value the utility of people living 500 years from now more or less as Nordhaus's DICE model would value them. Given the methodological constraint just mentioned, we seem barred from doing this unless we value an increase in the utility of one of our family members more than an equivalent increase in the utility of 1,710 currently existing strangers. Opinions may differ, but to me this would seem to require an excessive degree of clan loyalty.

4.2 Second limitation

Here is the second observation to which I want to draw our attention in this section. Discounting for kinship may justify us in caring more about those of our descendants who are nearer to us in time, since they are more likely to be more closely related to us, but it provides no justification for caring more about the welfare of unrelated strangers on the basis of their location in time.

Within the context of assessing optimal climate policy, this observation may prove especially significant. It is widely recognised that harms from climate change will fall principally on developing countries. These countries have the weakest obligations to pursue aggressive abatement policies, as recognised by their exemption from the group of Annex I countries bound by the Kyoto Protocol. The strongest obligations rest with developed countries. As Schelling (1995) and Lomborg (2001, 322) have emphasized, from the perspective of developed countries, a key question for the design of optimal climate policies is whether it is more desirable to benefit future inhabitants of developing countries by mitigating greenhouse gas emissions or to fund development programs that help poor inhabitants of developing countries in the here and now. To drive home the weight of the choice we face, Lomborg calculates that the cost to Annex I countries of complying with the Kyoto Protocol would be enough to fund the extension of basic health, education, water, and sanitation to every person in the developing world.

From the perspective of many current citizens of developed countries, both current and future inhabitants of developing countries are presumably equally unrelated. For such people, discounting for kinship would provide no basis for the choice of a positive rate of pure intergenerational time preference when deciding between abatement and development projects. A zero rate of pure intergenerational time preference is apparently the only appropriate option for addressing one of the most important prioritization decisions faced by these inhabitants of developed countries in reckoning with climate change.

4.3 Third limitation

Here is a third observation, which follows on naturally from the previous.

Under discounting for kinship, there apparently exists no unitary discount rate. Precisely because reasons for discounting are presumed to be agent-relative, there

will be no shared rate of pure time preference such that each currently existing individual ought to discount the welfare of future people at a certain rate per period, no matter which future people they might be. If reasons for pure intergenerational time preference are agent-relative, then we cannot ask whether and to what extent it may be justifiable to discount the interests of future people. We need to specify *for whom* it may be justifiable to discount the welfare of *which* people relative to *which other* people in light of their relative position in time.

This is not how discounting is standardly conceived by economists concerned with long-term policy setting, such as the analysis of optimal climate policy. There, it is standard to assume that there is some utility discount factor that is applied to future generations as a whole, whose rate of change is up for debate. This way of thinking seems to make little sense if we assume that reasons for pure intergenerational time preference are agent-relative, varying in character from person to person.

5 Global collectivism

The previous section outlined three important limitations of discounting for kinship. However, these limitations are not all that they appear to be. Contrary to the argument made in Sect. 4.3, in some contexts where economists have been especially concerned about the justifiability of pure time discounting, the question of what utility discount factor should be applied to future generations as a whole *is* appropriate under discounting for kinship. In these contexts, the assumed perspective is not that of some particular individual or nation, nor the wholly impartial point of view. It is something in-between. It is something like the point of view of *all of currently existing humanity*.

Section 5.1 introduces this idea, which I call *global collectivism*. Section 5.2 notes a series of important questions about the concrete interpretation of global collectivism in this context. Sections 5.3–5.5 discuss what we should make of the three observations noted in Sect. 4 in light of global collectivism, identifying ways in which the limitations on agent-relative pure time discounting previously noted drop away when viewed in this perspective.

5.1 Shared reasons of partiality

I suggested above that in some contexts where economists have been especially concerned about the rate of pure time preference, the assumed perspective may be that of the world as a whole. The analysis of optimal climate policy represents a clear example of this. As Stern (2008, 16) conceives of the issue, it concerns "social decisions by the world community now, bearing in mind consequences for future generations." Although he chides Stern for adopting "the lofty vantage point of the world social planner," (Nordhaus 2007, 691), Nordhaus (2008) adopts a similar approach. The social welfare function adopted in the DICE model is assumed to represent the collective preferences of the world as a whole. The optimal carbon tax recommended on the basis of the model is an internationally harmonized carbon

price "imposed in order to put the globe on the economically optimal path" (Nordhaus 2008, 196).

In this debate, therefore, the key question under discussion is what 'the world community now' should do, as opposed to what you or I should do, or what some particular country or bloc of countries should do (compare Nussbaum 2006, 279–81; Wringe 2005, 2014). Thus, insofar as there are reasons of partiality in play, those reasons may belong not to some particular individual, nor some particular country, but to this much greater collection of agents.

An analogy helps to clarify what I have in mind. Suppose that Kasei is Hiroko's son and Nikki is Nadezhda's daughter. Suppose these people are otherwise strangers to one another. Imagine that Kasei and Nikki are drowning, and so is a third person, Zoya. Zoya is a stranger to all of them. We stipulate that neither Hiroko nor Nadezhda can save any person on their own, but together they can save exactly two of the people who are drowning.

Intuitively, Hiroko and Nadezhda together have most reason to save Kasei and Nikki. Note, however, that Nadezhda has no reason to prefer that Kasei and Nikki are saved, as opposed to Nikki and Zoya. Kasei and Zoya are both equally strangers to Nadezhda. Similarly, Hiroko has no reason to prefer that Kasei and Nikki are saved, rather than Kasei and Zoya, since Nikki and Zoya are both equally strangers to Hiroko. Nonetheless, we think that Hiroko and Nadezhda together have most reason to save Kasei and Nikki. In this sense, there are reasons of partiality that they together have, which pick out the pair of Kasei and Nikki as uniquely important, but which no individual among them has.[15]

The suggestion I want us to consider, then, is that just as we intuitively describe Hiroko and Nadezhda as together having most reason to save Kasei and Nikki, so, in the same way, 'the world community now,' when engaged in internationally coordinated action in response to global climate change or other similar challenges, may be said to have reasons of partiality that belong to us collectively and which pick out the next generation as uniquely important, and subsequent generations as less so. 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.

5.2 Details to be decided

Call this way of understanding what discounting for kinship means in the context of problems requiring internationally coordinated action *global collectivism*. I have only given us a sketch of the idea. Many important details remain to be filled in.

Firstly, who exactly are the members of the 'the world community now'? Are we to think of this as the collection of all individual human beings currently living? Or

[15] Exactly how we should interpret the attribution of reasons of partiality to a group is discussed in Sect. 5.2. Note, in particular, that I do not assume that this attribution presupposes the existence of a collective agent, as opposed to being an instance of non-distributive plural predication. The latter interpretation strikes me as mandatory in this example.

should it instead be understood as a collection of states? In the context of global climate change, should it be thought of as the collection of all UNFCCC signatories?

Secondly, in asserting the existence of reasons that are not the reasons of some particular individual, but of a group of individuals, should we think of global collectivism as committing us to a view on which 'the world community now' is a collective agent who has these reasons? This might seem implausible if 'the world community now' is supposed to represent the collection of all currently existing individual human beings. This collection may seem to represent a so-called 'unstructured group,' lacking any shared procedure for collective decision-making, without which collective agency seems impossible (Collins 2019; French 1979, 1984, 1995; List and Pettit 2011, 158–159; Sheehy 2006). The attribution of collective agency is a lot more plausible if 'the world community now' is taken to refer to the UNFCCC signatories, since the High-Level Segment of the annual Conference of Parties may be thought to represent a procedure for collective decision-making of the kind we expect group agents to have.

Perhaps we need not suppose that collective reasons require collective agents. We could instead suppose that when we speak of reasons that are not the reasons of some particular individual, but of 'the world community now', this involves so-called *non-distributive plural predication* (McKay 2007; Oliver and Smiley 2013). In other words, we should think of ourselves as predicating something of a group of people that cannot be correctly predicated of any one among them, nor of a fusion of the group members (Björnson 2014; Pinkert 2014). Consider an analogy (McKay 2007, 24). It may be correct to say that the students in my class are seventeen in number. Clearly, no student in my class is seventeen in number. Nor is it plausible that an individual entity corresponding to a fusion of the students has the property of being seventeen in number. Any individual entity is only one in number. Thus, collectives can have properties had neither by their members nor by the collective considered as a supra-individual entity. Similarly, we may suppose that when we say that Hiroko and Nadezhda together have most reason to save Kasei and Nikki or that we together have greater reason to care about the next generation than about later generations, this involves non-distributive plural predication: attributing reasons of partiality to a group of people that we do not attribute to any one among them considered individually, nor to a supra-individual group agent constituted by those people.

For the present time, I elect to leave these issues of detail unanswered. In the remainder of this paper, I simply want to make the case for taking global collectivism seriously. To that end, I want us to consider what global collectivism tells us about the nature and significance of the three observations about the limitations of discounting for kinship noted in Sect. 4. Relying on the partial grasp that we currently have on the core idea, we are nonetheless able to see that there is a strong case for thinking of these limitations as being generally much less serious—or even non-existent—when viewed in light of global collectivism.

5.3 Third limitation revisited

I will work in reverse order, beginning with the third observation discussed in Sect. 4: namely, that if reasons for pure intergenerational time preference are agent-relative, then we cannot ask whether and to what extent it may be justifiable to discount the interests of future people *in general*, contrary to the practice of economists concerned with long-term policy setting like Nordhaus and Stern.

Global collectivism allows us to straightforwardly dismiss this concern. Under global collectivism, our reasons for caring differentially about the welfare of some people as opposed to others in virtue of their location in time are not understood as the reasons of some particular currently existing person or country. They are the reasons of the 'world community now'. More exactly, they are the reasons of the current generation, considered as a whole, to care about the welfare of each future generation, considered as a whole. Viewed from this perspective, there is nothing problematic about the standard assumption that there is a shared utility discount factor that may be applied to each future generation considered *en masse*, albeit one whose value declines as a function of time.

5.4 Second limitation revisited

Continuing our way back through Sect. 4, let's now consider the observation discussed in Sect. 4.2: namely, that discounting for kinship may justify us in caring more about those of our descendants who are nearer to us in time, but provides no justification for caring more about the welfare of unrelated strangers on the basis of their location in time. This limitation was suggested as being especially significant in the context of climate policy assessment.

Once again, this concern seems inapplicable in the context of global collectivism. From the perspective of 'the world community now,' there are no strangers, present or future. Speaking for the current generation as a whole, there are no human beings who will come to exist but who will not be our descendants. Therefore, under global collectivism, we need not worry that there exist certain groups of people relative to which we have no justification for caring more about those of its members who are nearer to us in time. The practice of pure time discounting will not be tightly circumscribed in the way suggested in Sect. 4.2.

5.5 First limitation revisited

Last but not least, consider the observation discussed in Sect. 4.1. There it was noted that if the application of a declining utility discount factor to the well-being of our descendants is justified in terms of discounting for kinship, this should never leave us in a position where we value the utility of one of us existing now over those of n distantly related descendants existing at time t unless we also value a current family member more than n currently existing strangers. It was suggested that this methodological prescription would call into question what might otherwise seem to be relatively modest suggestions for selecting a time schedule of utility discount factors.

When discounting for kinship is understood according to global collectivism, this prescription loses its bite. From the perspective of the world community as a whole, there presumably are no currently existing strangers. There is no one now living who is not a member of this community. Therefore, the constraint noted in Sect. 4.1 can be trivially satisfied by any schedule of discount factors.

However, we arguably should not read too much into this.[16] Even if there are no currently existing people who are not members of 'the world community now,' there presumably could have been. Exactly what this would mean will depend on exactly how we understand 'the world community now'. But however exactly we interpret this idea, it seems possible that there could now have existed human beings who aren't among its members: people whom we discovered living on other planets, say. Moreover, the fact that there are no such people is presumably morally arbitrary and should not change how steeply we are permitted to discount the welfare of future people.

We might, therefore, re-interpret our original constraint so that it now constrains how we trade off our own welfare against the welfare of our descendants in terms of how we would trade off our own welfare against the welfare of currently existing human beings whom we discovered living on other planets. Since it is morally arbitrary that no such people exist, this constraint should not be any easier to satisfy than the methodological constraint discussed in Sect. 4.1. We should therefore conclude that global collectivism does not, after all, permit the choice of a more extreme schedule of discount factors. As I see it, this is not a bug, but a feature.

5.6 Section summary

This section has outlined the idea of global collectivism as a framework for thinking about discounting for kinship in the context of assessing optimal climate policy or other similar problems requiring internationally coordinated action, such as the regulation of dangerous biotechnologies or risks from artificial intelligence. I have noted a number of important questions about how best to interpret global collectivism that are as yet unanswered. Even without resolving these questions, we are able to see that there is a strong case for thinking that global collectivism should lead us to think of the limitations on discounting for kinship highlighted in Sect. 4 as fading into insignificance in two out of three cases. Thinking of discounting for kinship in light of global collectivism allows us to answer the concerns noted in Sect. 4.2 and 4.3 in ways both straightforward and convincing. The story is more complicated for the observation made in Sect. 4.1. It may be true that there is no one now living who falls outside the scope of 'the world community now,' and so we can trivially satisfy the constraint that we should never end up in a position where we value the utility of one of us existing now over those of n distantly related descendants existing at time t unless we also value one of us existing now more than n currently existing strangers. However, since the absence of currently existing strangers is morally arbitrary, we are able to re-state the aforementioned

[16] I'm grateful to Alex Dietz for this observation.

methodological constraint in terms of counterfactual currently existing strangers so as to derive a no less restrictive constraint on the scheduling of discount factors.

6 Conclusion

Most moral philosophers think a positive rate of pure intergenerational time preference is unjustifiable. I have argued that it may be justified in terms of agent-relative reasons related to diminishing partiality between ever more distantly related generations, a view I call *discounting for kinship*. I have paid particular attention to what it means to discount the utility of future people in this way when our relationship to future people is viewed through the lens of (what I call) *global collectivism*, which involves attributing collective reasons of partiality to 'the world community now'. This view deserves further study. As I have noted, there are important questions about its interpretation and application that remain unanswered. I hope to have convinced you that these questions are worth answering, because the package of discounting for kinship and global collectivism represents an especially promising framework for thinking about the ethics of discounting in cases requiring internationally coordinated action, such as the questions concerning optimal climate policy over which economists have clashed. I have given us only a sketch of the framework. I hope this paper serves as encouragement for us to complete the sketch, and I hope to take up this project in future work.

In closing, I also want to remind us that the points developed in Sect. 5 can arguably be transposed with minimal modification to inform other ways in which we may choose to develop the idea that a positive rate of pure intergenerational time preference can be justified in terms of special obligations whose strength weakens over time, such as the communitarian approach noted in Sect. 3.1. While I continue to maintain that discounting for kinship is the most straightforward and intuitively compelling development of the core idea, the dialectical moves explored in this paper are arguably available more broadly to those who wish to explore alternative interpretations of the proposal that special obligations can justify a positive pure rate of time preference.

Acknowledgements For helpful comments on previous drafts of this paper and/or its core ideas, I would like to thank Adam Bales, John Broome, Alex Dietz, Mattias Endres, Hilary Greaves, Harry Lloyd, Sven Neth, Christian Tarsney, Teru Thomas, David Thorstad, and Phil Trammel, as well as audiences at the Interdisciplinary Research Network for Economists and Philosophers in York in September 2019 and at the 2nd Oxford Workshop on Global Priorities Research in Oxford in July 2019. I would also like to thank the anonymous referees who have reviewed the paper and offered helpful comments and criticisms.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

Ahmed, A. (2020). Rationality and future discounting. *Topoi, 39*, 245–256.

Arrow, K. J. (1996) Discounting, morality, and gaming. Working Paper 97004, Stanford University, Department of Economics.

Beckerman, W., & Hepburn, C. (2007). Ethics of the discount rate in the stern review on the economics of climate change. *World Economics, 8*(1), 187–210.

Björnsson, G. (2014). Essentially shared obligations. *Midwest Studies in Philosophy, 38*(1), 103–120.

Boonin, D. (2003). *A defense of abortion*. Cambridge University Press.

Boonin, D. (2014). *The non-identity problem and the ethics of future people*. Oxford University Press.

Broome, J. (1992). Counting the cost of global warming. White Horse Press.

Broome, J. (1994). Discounting the future. *Philosophy and Public Affairs, 23*(2), 128–156.

Broome, J. (2004). *Weighing lives*. Oxford University Press.

Broome, J. (2012). *Climate matters: ethics in a warming world*. NW. W. Norton & Company.

Broome, J. (2016). The well-being of future generations. In M. D. Adler & M. Fleurbaey (Eds.), *The Oxford handbook of well-being and public policy* (pp. 901–928). Oxford University Press.

Caney, S. (2009). Climate change and the future: Discounting for time, wealth, and risk. *Journal of Social Philosophy, 40*(2), 163–186.

Caney, S. (2014). Climate change, intergenerational equity and the social discount rate. *Politics, Philosophy & Economics, 13*(4), 320–342.

Collins, S. (2019). *Group duties: Their existence and their implications for individuals*. Oxford University Press.

Cowen, T., & Parfit, D. (1992). Against the social discount rate. In J. S. Fishkin & P. Laslett (Eds.), *Justice between age groups and generations* (pp. 144–161). Yale University Press.

Dasgupta, P. (2008). Discounting climate change. *Journal of Risk and Uncertainty, 37*, 141–169.

De-Shalit, A. (1995). *Why posterity matters: Environmental policies and future generations*. Routledge.

Foot, P. (1967). The problem of abortion and the doctrine of double effect. *Oxford Review, 5*, 5–15.

French, P. A. (1979). The corporation as a moral person. *American Philosophical Quarterly, 16*(3), 207–215.

French, P. A. (1984). *Collective and corporate responsibility*. Columbia University Press.

French, P. A. (1995). *Corporate ethics*. Harcourt Brace.

Gollier, C. (2013). *Pricing the future: The economics of discounting and sustainable development*. Princeton University Press.

Greaves, H. (2017). Discounting for public policy: A survey. *Economics and Philosophy, 33*(3), 391–439.

Hanser, M. (2008). The metaphysics of harm. *Philosophy and Phenomenological Research, 77*(2), 421–450.

Hanser, M. (2009). Harming and procreating. In D. Wasserman & M. A. Roberts (Eds.), *Harming future persons* (pp. 179–199). Springer.

Harman, E. (2004). Can we harm and benefit in creating? *Philosophical Perspectives, 18*, 89–113.

Harman, E. (2009). Harming as causing harm. In D. Wasserman & M. A. Roberts (Eds.), *Harming future persons* (pp. 137–154). Springer.

Heal, G. M. (1998). *Valuing the future: Economic theory and sustainability*. Columbia University Press.

Henderson, N., & Bateman, I. (1995). Empirical and public choice evidence for hyperbolic social discount rates and the implications for intergenerational discounting. *Environmental and Resource Economics, 5*(4), 413–423.

Hepburn, C. (2006). Discounting climate change damages: working note for the Stern review. LSE Working Paper. Retrieved January 4 2022 at http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.129.8102&rep=rep1&type=pdf

Hurka, T. (2005). Proportionality in the morality of war. *Philosophy and Public Affairs, 33*(1), 34–66.

Intergovernmental Panel on Climate Change (2015). Social, economic, and ethical concepts and methods. In Climate change 2014: mitigation of climate change: working group III contribution to the IPPC fifth assessment report (pp. 207–282). Cambridge University Press.

Jackson, F. (1997). Which effects? In J. Dancy (Ed.), *Reading Parfit* (pp. 42–53). Blackwell.

Jamieson, D. (2014). *Reason in a dark time: Why the struggle against climate change has failed – and what it means for our future*. Oxford University Press.

Kamm, F. M. (2004). Failures of just war theory: Terror, harm, and justice. *Ethics, 114*(4), 650–692.

Kelleher, J. P. (2017). Pure time preference in intertemporal welfare economics. *Economics and Philosophy, 33*(3), 441–473.
Keller, S. (2013). *Partiality*. Princeton University Press.
Kolodny, N. (2010). Which relationships justify partiality? The case of parents and children. *Philosophy and Public Affairs, 38*(1), 37–75.
Lackey, D. P. (1984). Moral principles and nuclear weapons. Rowman & Allanheld.
Lefkowitz, D. (2009). Partiality and weighing harm to non-combatants. *Journal of Moral Philosophy, 6*(3), 298–316.
List, C., & Pettit, P. (2011). *Group agency: The possibility, design, and status of corporate agents*. Oxford University Press.
Lomborg, B. (2001). *The skeptical environmentalist: Measuring the real state of the world*. Cambridge University Press.
Lord, E. (2016). Justifying partiality. *Ethical Theory and Moral Practice, 19*(3), 569–590.
MacIntyre, A. (1984). *Is patriotism a virtue? The Lindley Lecture*. University of Kansas Press.
McKay, T. J. (2007). *Plural predication*. Oxford University Press.
McMahan, J. (2002). *The ethics of killing: Problems at the margins of life*. Oxford University Press.
Nagel, T. (1970). *The possibility of altruism*. Oxford University Press.
Norcross, A. (2005). Harming in context. *Philosophical Studies, 123*(1/2), 149–173.
Nordhaus, W. D. (2007). A review of the stern review on the economics of climate change. *Journal of Economic Literature, 45*(3), 686–702.
Nordhaus, W. D. (2008). *A question of balance: Weighing the options on global warming policies*. Yale University Press.
Nussbaum, M. C. (2006). *Frontiers of justice: Disability, nationality, species membership*. Harvard University Press.
Oliver, A., & Smiley, T. (2013). *Plural logic*. Oxford University Press.
Parfit, D. (1984). *Reasons and persons*. Oxford University Press.
Parfit, D. (2011). *On what matters:* (Vol. 1). Oxford University Press.
Pinkert, F. (2014). What we together (can be required) to do. *Midwest Studies in Philosophy, 38*(1), 187–202.
Ramsey, F. P. (1928). A mathematical theory of saving. *The Economic Journal, 38*(152), 543–559.
Rawls, J. (1999). *A theory of justice* (revised). Harvard University Press.
Rendall, M. (2019). Discounting, climate change, and the ecological fallacy. *Ethics, 129*(3), 441–463.
Scheffler, S. (1997). Relationships and responsibilities. *Philosophy and Public Affairs, 26*(3), 189–209.
Scheffler, S. (1999). Conceptions of cosmopolitanism. *Utilitas, 11*(3), 255–276.
Scheffler, S. (2004). Projects, relationships, and reasons. In R. J. Wallace, P. Pettit, S. Scheffler, & M. Smith (Eds.), *Reason and value: Themes from the moral philosophy of Joseph Raz* (pp. 247–269). Oxford University Press.
Schelling, T. (1995). Intergenerational discounting. *Energy Policy, 23*(4–5), 395–401.
Sheehy, P. (2006). Holding them responsible. *Midwest Studies in Philosophy, 30*(1), 74–93.
Shiffrin, S. V. (1999). Wrongful life, procreative responsibility, and the significance of harm. *Legal Theory, 5*(2), 117–148.
Stern, N. (2006). *The economics of climate change: The Stern review*. Cambridge University Press.
Stern, N. (2008). The economics of climate change. *American Economic Review, 98*(2), 1–37.
Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *The Review of Economics Studies, 23*(3), 165–180.
Thomson, J. J. (1985). The trolley problem. *Yale Law Journal, 94*(6), 1395–1415.
Velleman, J. D. (2005). Family history. *Philosophical Papers, 34*(3), 357–378.
Velleman, J. D. (2008). Persons in prospect. *Philosophy and Public Affairs, 36*(3), 221–288.
Voorhoeve, A. (2014). How should we aggregate competing claims? *Ethics, 125*(1), 64–87.
Williams, B. (1982). Persons, character, and morality. In his *Moral luck: Philosophical papers 1973–1980* (pp. 1–19). Cambridge University Press.
Wolf, S. (1992). Morality and partiality. *Philosophical Perspectives, 6*, 243–259.
Woollard, F. (2012). Have we solved the non-identity problem? *Ethical Theory and Moral Practice, 15*(5), 677–690.
Wringe, B. (2005). Needs, rights, and collective obligations. *Royal Institute of Philosophy Supplement, 57*, 187–208.

Wringe, B. (2014). From global collective obligations to institutional obligations. *Midwest Studies in Philosophy, 38*(1), 171–186.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.