

Time Bias and Altruism

Leora Urim Sung (University College London)

Global Priorities Institute | December 2023

GPI Working Paper No. 17-2023

Please cite this working paper as: Sung, L. U. (2023) Time Bias and Altruism
Global Priorities Institute Working Paper Series, No. 17-2023 Available at:
<https://globalprioritiesinstitute.org/leora-urim-sung-time-bias-and-altruism>



Time Bias and Altruism

Leora Urim Sung¹

Abstract

We are typically *near-future biased*, being more concerned with our near future than our distant future. This near-future bias can be directed at others too, being more concerned with their near future than their distant future. In this paper, I argue that, because we discount the future in this way, beyond a certain point in time, we morally ought to be more concerned with the present well-being of others than with the well-being of our distant future selves. It follows that we morally ought to sacrifice our distant-future well-being in order to relieve the present suffering of others. I argue that this observation is particularly relevant for the ethics of charitable giving, as the decision to give to charity usually means a reduction in our distant-future well-being rather than our immediate well-being.

1. Introduction

People are typically *time biased* with respect to their well-being. For instance, we often display *future bias*, being more concerned with our future well-being than with our past well-being. In addition to future bias, many people also display *near-future bias*, being more concerned with their near-future well-being than with their distant-future well-being.

In this paper, I will argue that, because we display near-future bias, if we care enough about other people, there will be a point in time at which we care more about the present condition of other people than our distant-future condition. And since

¹ For helpful comments and insightful conversations, I thank Joe Horton, Kacper Kowalczyk, Christian Tarsney, Teruji Thomas, Nikhil Venkatesh, and participants of the 2022 Early Careers Conference Programme (Global Priorities Institute, University of Oxford).

we morally ought to have a sufficient level of concern for other people, it follows that we morally ought to care more about the present condition of other people than about our distant-future condition. If this is right, then it also follows that we morally ought to sacrifice our distant-future well-being in order to relieve the current suffering of others. I outline this argument in Section Three, after giving a more thorough explanation of self-regarding and other-regarding time bias in Section Two.

In Section Four, I draw out a practical implication of this observation. The claim that we ought to sacrifice our distant-future well-being to relieve the current suffering of others is particularly relevant for the ethics of charitable giving. This is because the decision to give to charity usually leads not to a reduction in the agent's immediate well-being but rather to a reduction in the agent's distant-future well-being. So, my argument calls into question the morality of saving up to secure our distant future when there are currently millions of people starving across the world.

In Section Five, I look at three types of potential objection to my argument: First, I address the objection that we are not morally obligated to be so concerned with the welfare of distant strangers as to require us to be more concerned for their well-being than our distant-future well-being or the well-being of our loved ones. Second, I address the objection that moral agents rationally ought to be temporally neutral rather than display time bias. Third, I address worries relating to economic facts about the world which may seem to undermine my argument.

2. Self-Regarding and Other-Regarding Time Bias

There are two ways in which most of us fail to be temporally neutral with respect to our own interests. In this section, I explain both, and I argue, following Caspar Hare, that we ought to be correspondingly temporally a-neutral towards the interests of other people.

People often display future bias, showing asymmetrical attitudes toward future and past pains and pleasures. We strongly prefer painful experiences to be in the past and pleasurable experiences to be in the future. And not only do we prefer painful experiences to be in the past, but we would also even prefer to have experienced a longer period of pain in the past than to experience a shorter period of pain in the

future. Derek Parfit illustrates this aspect of our future bias with the following thought experiment:

My Past or Future Operation. You are in the hospital to have an extremely painful but safe operation for which you can be given no anaesthetic. In order to ease recovery, you know that the hospital will give you drugs that cause you to forget the operation as soon as it is completed. You wake up in hospital and the nurse informs you that either you *have* undergone a painful four-hour operation, or you *will soon* undergo a painful one-hour operation.²

Most people would prefer to have already had the four-hour operation than to face the one-hour operation in the future. They would be immensely relieved if it turns out that their ordeal is over. This shows that we not only prefer painful experiences to be in the past, but also prefer that our lives contain more total hours of pain, if it means less of the pain is still to come.

When it comes to pleasurable experiences, our preferences for whether we want the event to be in the past or future switch. We prefer pleasurable experiences to be in the future rather than the past, and we would even prefer to experience a shorter period of pleasure in the future than to have experienced a longer period of pleasure in the past. In other words, we prefer our lives to contain fewer total hours of pleasure if this means more of the pleasure is still to come.

In addition to future bias, many people also have a bias toward the near future. We care less about what is good or bad for us the further into the future it will happen, preferring pleasurable experiences to be in the near future and painful experiences to be in the distant future. Sometimes, we would prefer our lives to contain fewer total hours of pleasure if that means we experience the pleasure in the near future rather than the distant future. For example, imagine you are faced with the choice of going on a pleasant two-day skiing vacation this year, or a pleasant three-day skiing vacation in five years. I suspect many people would choose the former over the latter, all other things being equal. We might also prefer our lives to contain more total hours of pain if that means we experience the pain in the distant future rather than in the immediate

² See Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press 1984): 165.

future. For example, we often put off dental treatments even though it means more suffering in the long run. We tend to avoid pain in the immediate future at the cost of a worse experience in the distant future. So, it seems we are naturally inclined to discount the distant future, both with regards to pleasurable experiences and painful experiences.

Our preferences might be less clear when we consider what we want for other people. Philosophers have pointed out that our tendencies to display time bias typically disappear when it comes to other people's interests. For instance, Parfit describes there being 'a surprising asymmetry in our concern for our own, and other people's pasts'.³ While we feel relieved knowing that some ordeal of ours is over, we do not experience the same relief knowing that some ordeal has already occurred to a loved one. Consider the following alteration to Parfit's thought experiment:

My Daughter's Future or Past Operation. Your daughter has just woken up in hospital. The nurse informs you that there has been a mix up with the patient charts. Your daughter has either already had a painful four-hour operation or is about to undergo a painful one-hour operation.

Should you prefer that she has just had a four-hour operation, or that she will soon undergo a one-hour operation? It is not so clear. If you prefer that she has a better overall *life*, a life in which she experiences fewer hours of pain, you should prefer that she will undergo a one-hour operation in the future. If you prefer that she be in a better *predicament*, you should prefer that she has already had a four-hour operation in the past.

Caspar Hare thinks that most people's preferences will switch depending on where their daughter is.⁴ He suggests that if she is far away on the other side of the world, we will be temporally impartial and prefer that she has a better overall life—that is, we will prefer that she is to undergo the one-hour operation in the future rather than that she is recovering from the four-hour operation. If your daughter is close by,

³ Parfit (1984): 182

⁴ Caspar Hare, 'A Puzzle About Other-Directed Time Bias', *Australasian Journal of Philosophy* 86 (2008): 269- 277.

lying in the hospital bed in front of your eyes, Hare suggests we will display future bias and care more for her predicament, preferring that she has already gone through her ordeal.

Hare argues that there are no good reasons to think that appropriate concern for one's daughter should mandate this switch in preferences. You should prefer, in both cases, either for her to have had the longer operation in the past, or that she will have a shorter operation in the future. And Hare argues for the latter preference as follows. Since most people are future-biased, your daughter will almost certainly prefer that she be in the better predicament. And it seems strange that the preferences mandated by your appropriate concern for her would contradict her preferences. So, it seems that, in both cases, you ought to display future bias on her behalf, and prefer that she has already had the longer operation. Hare concludes that the temporal impartiality induced in us by distance from the objects of our concern is the result of our failing to engage, imaginatively, with their present condition.

Hare's claims can be supported by empirical evidence. There are conflicting results regarding whether we display other-regarding time bias. In a study conducted by Caruso, Gilbert and Wilson, most of the participants displayed other-regarding time-neutrality, considering all parts of the other person's life with equal consideration.⁵ However, in Greene et al.'s study, most of the participants wanted good experiences to be in the other person's future, and bad experiences to be in the other person's past.⁶ Greene et al. offer the following hypothesis to explain this contradiction in findings. While the participants in the Caruso et al. study were prompted to consider the fate of an unidentified third party, in the Greene et al. study, participants were given rich biographical information about the other, such that they identified with the third party and thereby adopted their preferences. This supports Hare's claim that other-regarding time-neutrality is a result of being detached from the other person and failing to engage with their condition. If we know enough about the third party, we start to identify with them and begin to adopt their preferences, displaying other-regarding future bias.

⁵ Eugene Caruso, Daniel Gilbert, and Timothy Wilson. "A Wrinkle in Time: Asymmetric Valuation of Past and Future Events" *Philosophical Science* 19/8 (2008): 796-801.

⁶ Preston Greene, Andrew Latham, Kristie Miller, and James Norton. "Hedonic and Non-Hedonic Bias Toward the Future" *Australasian Journal of Philosophy* 99 (2021): 148-63.

In the same way that we display future bias toward other people we identify with, our near-future bias can also be directed at others too. If Hare is right, just as we care more about our present- and immediate-future interests than our distant-future interests, appropriate concern for other people should mandate that we care more about their present- and near-future interests than their distant-future interests. This is because it would be strange that the preferences mandated by our appropriate concern for others would contradict their preferences. Not only should we display other-regarding future bias, but we should display other-regarding near-future bias as well.

This near-future bias directed at others means that we should be more concerned with alleviating their present suffering than with alleviating their distant-future suffering. If we remain temporally neutral, and we are indifferent about whether they are suffering now or whether they will suffer some time in the future, this seems to show a failure to engage, imaginatively, with their present condition. When we do imaginatively engage with their present condition, there is a greater immediacy or urgency with the present suffering of people than learning of some misfortune that will come upon them many years in the future. In other words, we feel a stronger compulsion to alleviate their present- and near-future suffering over the suffering they will experience in the far-distant future.

In the next section, I will explain the implications of our self-regarding and other-regarding near-future bias for how we should weigh other people's present interests against our distant-future interests.


3. Near-Future Bias and Concern for the Other

If I am impartial, or agent neutral, I will give the same weight to the welfare of other people as to my own welfare. This means that the present interests of other people will matter to me just as much as my own present interests. The future interests of other people will also matter just as much as my own future interests. Given my near-future bias, it follows that I will always be more concerned with other people's present interests than with my own future interests.

Sometimes, out of love, I might go beyond agent neutrality and give *more* weight to the interests of other people than to my own interests. In such cases, the present interests of this loved one will matter more to me than my own present interests. I will also care about their future interests more than I care about my own future interests. I might even care more about their *future* condition than my own *present* condition, despite displaying near-future bias both regarding my own well-being and that of other people. We often see this with parents who sacrifice their present well-being in order to secure their children's future well-being.

When it comes to distant strangers, however, we would not expect the agent to show such sacrificial love or even agent-neutrality. Most people would accept that it is permissible for me to give greater weight to my present and future well-being, than to the present and future well-being of distant strangers. However, given my near-future bias, I will care less about my distant future than my present. And, as I argued in the previous section, if I show appropriate concern for others by taking on their preferences, I will show the same near-future bias regarding their present well-being and their future well-being. In other words, I will care more about their present condition than about their distant future condition.

It follows that, if my level of concern for the interest of others exceeds a particular threshold, there will be a point at which I will care more about other people's present well-being than about my distant-future well-being. The following graph illustrates all this:



This graph shows that, even if I grant greater weight to my own interests than to the interests of strangers, so long as my concern for the stranger is above a certain threshold, there will be a future point in time, T_n , such that my concern for my well-being after this point will be equal to or less than my concern for the stranger's current well-being. I will be indifferent between giving an extra unit of well-being to myself at T_n and giving the same unit to the distant stranger now. And my concern for my well-being beyond T_n will, in fact, be less than my concern for the current well-being of the stranger. In other words, I will prefer to provide an extra unit of well-being to the stranger now than to my distant-future self at any time beyond T_n . Since I morally ought to have a decent amount of concern for others, it seems to follow that, given my near-future bias, I morally ought to sacrifice my well-being beyond T_n for the sake of the current well-being of distant strangers.

Of course, whether there is such a point T_n will depend on the position of the threshold for a 'decent amount of concern' for others, and how much I discount my distant-future interests. It could be that there is no point at which I am equally concerned with a stranger's present interests over my distant-future interests, either because the threshold is too low or because I barely discount my distant-future interests. I will return to this worry in Section Five. But for the time being, I will assume that the rate at which I discount my distant-future and the level of concern I ought to have for distant strangers is such that there *is* a point at which I morally ought to be more concerned with the present interests of strangers than with my distant-future interests, and therefore I am morally obligated to sacrifice my distant-future well-being for the sake of the current well-being of a stranger.

In the next section, I argue that these observations have important implications for the ethics of charitable giving and our moral obligations to alleviate poverty.

4. Near-Future Bias and Altruism

I have argued that, given our near-future bias regarding both our own interests and the interests of other people, if we are sufficiently concerned for others, there will be a point at which we care more about other people's present interests than about our own distant-future interests. It follows that, beyond this point in time, we morally

ought to sacrifice our distant-future interests for the sake of other people's present interests.

This observation has particular importance for the ethics of charitable giving. This is because, when it comes to beneficence, it is usually not the agent's present well-being that is placed in competition with other people's present well-being. Giving to charity does not usually require the agent to sacrifice their present well-being to alleviate the suffering of others. To see this, consider the following example:

Cinema. I am considering watching a movie at the cinema. I realise, however, that I could do a lot of good by spending my cinema money on famine relief, and so I give the money to a charity instead. The next day, I am considering watching the movie again, but I am faced with the same choice—once again, I could do a lot of good by spending my cinema money on famine relief.

Philosophers who argue that we have strong moral obligations to give to charity would say that each time I face this choice, I am morally obligated to give to charity rather than go to the cinema.⁷ Most people, however, do not believe our duties to alleviate poverty are so strong. So, in most real-life versions of this example, in most developed countries, the charitable agent is likely to *both* donate to charity *and also* go to watch the movie. That is, rather than sacrificing the trip to the cinema, the agent is likely to sacrifice a part of their savings.

There are, of course, people who cut back on cups of coffee, going to the movies, and other luxuries in order to give to charity. But assuming that we are talking about people of a certain level of financial security, the decision to give to charity does not usually create immediate suffering and loss for the agent. The agent is most likely to enjoy these luxuries in life while also donating to charity. For many people, charitable giving does not entail that we sacrifice our current well-being for the sake of a distant stranger's current well-being.

The decision to donate, however, while not creating immediate loss to the agent's well-being, will most likely affect her distant-future well-being in some way.

⁷ See Peter Singer, "Famine, Affluence, and Morality" (1977) and Peter Unger, *Living High and Letting Die*.

This is because a lifestyle of giving added up over the years will mean that the agent's distant-future well-being is compromised to a certain extent. For instance, instead of giving to charity, perhaps I could have invested the extra money, potentially greatly increasing my wealth many years in the future. The monthly donations added up over my lifetime may mean that I put less into my savings account, preventing me from enjoying a cushy retirement in my old age. So, unless I believe we have strong duties of beneficence, and assuming that I am of a certain level of financial security, while the decision to give to charity will not affect my present and immediate future well-being, it will most likely affect my distant future well-being in a significant way. So, for any moral agent over a certain level of financial security, the decision to give to charity is better framed as a trade-off between increasing a stranger's present well-being at some cost to their own distant-future well-being.

Whether we should give to charity, then, seems related to how much weight we should give to our future selves over the present suffering of others. I argued above that we are biased toward our near future, and that appropriate concern for others mandates that this bias is directed at others too. Given that we discount the value of our distant-future well-being, if we care enough about the well-being of others, there will be a point at which we will care more about the present well-being of others than about our distant-future well-being. So, if we are near-future biased, we morally ought to sacrifice our distant-future well-being in order to relieve the current suffering of others. This is the case even if we do not have strong obligations to alleviate poverty, and even if we are permitted to grant much more weight to our own interests than to the interests of others.

This calls into question the morality of saving up to secure our future when there are currently millions of people suffering around the world. Of course, this isn't to say that we should save nothing for our retirement, but rather to say that our savings policy should reflect our near-future bias. Given that we discount our future well-being the way we do, if we have the level of concern for the well-being of distant strangers that we should, we morally ought to be more concerned about alleviating their present suffering than securing some extra unit of well-being for our distant-future self. If this is right, it follows that we morally ought to be directing our extra

financial resources towards alleviating the present suffering of others rather than saving up for a cushy retirement or heavily investing in our distant future.⁸

5. Objections

There are several ways we might push back against, or object to, the argument I make above. First, we might dispute the level of concern we ought to have for distant strangers over our own interests and the interests of our loved ones. We might argue that we are not obligated to be so concerned with the well-being of distant strangers that we are morally required to sacrifice our distant-future well-being for the sake of their current well-being. Or we might argue that while it is not permissible for me to put my own interests way above that of strangers, we are permitted to do this with the interests of our loved ones, in a way that makes charitable giving remain non-obligatory. Second, we might argue that near-future bias is irrational, so that we should not display near-future bias, either regarding our own future interests or the future interests of others. Third, we might appeal to economic facts about the world that seem to undermine my argument. In this section, I will address each type of objection in turn.

5.1. The Morally Required Level of Concern


I claimed above that so long as I have a decent amount of concern for others, and given that I am near-future biased, there will be a point in time, T_n , such that my concern for my well-being at this point will be equal to my concern for a stranger's current well-being. However, whether there is such a point T_n will depend on the position of the threshold for a 'decent amount of concern' for others. In this subsection, I look at two ways in which we can challenge my claim that the level of concern I morally ought

⁸ This gives us reason to prioritise charities which focus on alleviating the beneficiary's present suffering, rather than charities where the payoff is in the beneficiary's distant future, such as charities which focus on education. However, if, for example, the returns of investing in education are sufficiently great, then we could justify donating to charities which focus on such long-term goals. I thank Nikhil Venkatesh for pressing me on this point.

to have for distant strangers relative to my own well-being is such that we ought to sacrifice our distant-future interests for the present interests of others.

5.1.1. *Less Concern for Distant Strangers*

One could argue that there is never a point at which I morally ought to care more about a distant stranger's present well-being than my distant-future well-being because we are just not required to give that much weight to the interests of distant strangers relative to our own interests. If we are permitted to give less weight to the interests of strangers than I have suggested we ought to give, the graph comparing the level of concern for my own welfare and that of a stranger over time would instead look something like this:




Here, we can see that even if I am near-future biased, there is never a point in time at which my concern for my own well-being is equal to or more than that of a stranger's well-being. I am permitted to prefer my own well-being, at any point in my life, to the well-being of strangers, at any point in their lives. It follows that I am not morally required to sacrifice my distant-future well-being for the sake of a stranger's present well-being.

However, even if we accept that there is never a point in time at which our concern for our own well-being is *equal* to that of a stranger's well-being, there is an

upper limit to how much we can put our own interests above the interests of others. Although a reasonable level of personal partiality can be permitted, it seems that the gap between the level of concern I have for my own well-being and that of strangers should not be exceedingly big. It should not, for instance, be so great as to permit me to walk away from a stranger before me in dire need, when I could help that stranger at very little cost to myself. So, there is a maximally permissible gap between the level of concern I can have for my present well-being and the present well-being of a stranger.

What my observation about our near-future bias shows, then, is that the maximally permissible gap between the level of concern I ought to have for my distant-future well-being and the present well-being of strangers is significantly less than the gap for my present well-being and the present well-being of strangers. Whatever you think is the maximum permissible weight you can give to your present well-being relative to the present well-being of strangers, the maximum weight you can permissibly give to your distant-future well-being relative to the present well-being of strangers will be significantly less, given that you discount your future well-being. So, even if there is never a point in time at which I should be *more* concerned about the present condition of strangers than about my distant-future condition, I still ought to be *significantly less* concerned about my distant-future condition relative to the present condition of strangers. The following graph illustrates this:



Distance a is the gap between the maximum level of concern I am permitted to have for my present interests over the present interest of a stranger. Distance b is the gap between the level of concern I have for my distant-future and the level of concern I have for the present interest of a stranger. As we can see, distance b is smaller than distance a . This implies that, even if I am not required to be equally concerned for a distant stranger's well-being over my well-being at any given time, I should at least be willing to sacrifice some unit of my distant-future well-being in order to greatly increase a stranger's present well-being.

It seems that this is precisely the situation that faces us when it comes to alleviating poverty, for two reasons.

First, the diminishing marginal utility of wealth means that each incremental increase in wealth provides a smaller incremental increase in utility. In other words, as someone's income increases, they will receive a correspondingly smaller increase in satisfaction and happiness. This means that the extent to which we can increase someone's well-being for the same amount of money will be different depending on their financial situation. If I am reasonably well off, I can do a lot more good by donating a given amount of money to those in absolute poverty than I can by spending that money on myself. That is, I can greatly increase a poorer person's well-being at a comparatively small sacrifice to myself. For instance, if I spend an extra \$100 on myself this month, perhaps I can increase my well-being slightly by eating out a few more times or getting a few Ubers instead of the bus. The same \$100 is what it costs to provide Ready-to-use Therapeutic Food (RUTF) over a six to eight-week period to save a child suffering from severe acute malnutrition.⁹

Second, as I have observed, for those above a certain level of financial security, giving to charity will likely involve sacrificing their distant-future well-being rather than their current well-being. And given that we discount our future well-being because of our near-future bias, when we give to charity, we are in fact giving up a much smaller weighted unit of our well-being in order to greatly increase a stranger's well-being. For instance, say that I decide to donate \$1,200 over a period of a year in order to provide life-saving medicine, food and treatment for a malnourished child for

⁹ United Nations Children's Fund, *Severe wasting: An overlooked child survival emergency*, UNICEF, New York, May 2022, p. 8. <https://www.unicef.org/child-alert/severe-wasting>.

a year. If I am pretty well off anyway, I will probably not change my current lifestyle in order to donate this sum of money, but rather, I will put less money into my savings account. So, I am sacrificing something I regard to be of less value than \$1,200, considering that it will not affect my immediate well-being.

So, even if we are permitted to give much more weight to our own interests, the diminishing marginal utility of wealth coupled with the fact that we are forfeiting our far-distant future well-being rather than our immediate well-being makes it hard to deny that we have at least *some* obligations to alleviate the suffering of distant strangers rather than investing money in our own future. The upshot is that we probably should not be valuing our own interests so much as to save up for a cushy retirement at the expense of the immediate suffering of others.

5.1.2. *More Concern for My Loved Ones*

One might also argue that, while we are not permitted to grant our own interests too much weight over that of strangers, we are permitted to put the interests of our loved ones way above that of distant strangers. We may regard giving too much weight to our own well-being to be egoistical and narcissistic, but a great deal of concern for the well-being of a loved one is usually seen as a virtue, rather than a vice. Indeed, we consider parents who sacrifice their well-being for the sake of their children to be displaying characteristics that are admirable. So, it seems that the maximally permissible gap between the level of concern I have for someone I love and the level of concern for a stranger can be much greater than the gap between the level of concern I have for my own well-being and that of a stranger.

If this is true, while I may be morally required to sacrifice some unit of my distant-future well-being to alleviate a stranger's present suffering, I would not be morally required to sacrifice the distant-future well-being of my child or someone I love in order to alleviate a stranger's present suffering. This means that, while I may not be permitted to secure my own distant-future well-being at the expense of a stranger's current well-being, I would be permitted to secure the distant-future well-being of my child or someone I love. For instance, suppose I am right that appropriate concern for the well-being of strangers means that I am not permitted to save up for a cushy retirement rather than donate a certain proportion of that money to charity. If

I am permitted to grant much greater weight to the interests of my loved ones than I am permitted to grant to my own interests, it seems I could still be permitted to put that money into my child's savings account or keep it aside to leave them a hefty inheritance.

My response to this argument is that, although I do think we are permitted to grant even more weight to the interests of our loved ones than we are permitted to grant to ourselves, the additional weight is not unbounded, and it is not great enough to do the work the objection requires it to do. While it is admirable for a parent to sacrifice their present well-being for their child's near future, it seems too much if they do so for the child's far-distant future. If a parent is so concerned for the welfare of their child that they sacrifice their present well-being in order to secure their child's pension, we would likely think this is a case of overparenting and that there is something unhealthy about the relationship.¹⁰ So, while it does seem permissible to grant more weight to the interests of our loved ones than to ourselves, there seems to be a limit to the level of concern we morally ought to have for them over the interests of strangers.

So, there is a maximally permissible gap between the level of concern I can have for my child's present well-being and the present well-being of a stranger. Again, my observation about our near-future bias shows us that the maximally permissible gap between the level of concern I can have for my child's distant future well-being and the present well-being of strangers should be significantly less than that. Even if I am permitted to grant greater weight to my child over strangers, and even over my own interests, I should probably not be so concerned with their interests as to save up for my child's pension or leave them with a hefty inheritance at the expense of the immediate suffering of others. This point becomes even stronger when we consider that the diminishing marginal utility of wealth will also apply when it comes to what we can provide for our children in comparison to what we can provide for distant strangers in absolute poverty.

¹⁰ There are, of course, exceptions to this. For instance, I may have a child who will, through some disability, never become an independent adult. We would not consider it to be overparenting to ensure that the distant future of such a child is provided for. I thank Teruji Thomas for pressing me on this.

5.2. The Temporally Neutral Agent

While many people accept the rationality of *future* bias, the consensus among philosophers seems to be that *near-future* bias is a rational defect.¹¹ In other words, the consensus among philosophers seems to be that it is irrational to be more concerned about our near-future well-being than our distant-future well-being. This is because, the reasoning goes, the rational person would make choices that result in their leading better lives. If this is true, it seems we should not display near-future bias, both with regards to our own future interests and with regards to the interests of others. So, we should not accept the observation I made above about putting the present well-being of strangers above our distant-future well-being.

As I will now explain, there are several ways to push back against this line of objection.

5.2.1. *The Irrelevance of the Irrationality of Near-Future Bias*

It is important to note that my argument does not necessarily depend on the *rationality* of near-future bias, but rather, it is grounded on the fact that people *do* display this kind of bias. For given that we display near-future bias, appropriate concern for other people mandates that we take on their preferences and display other-regarding near-future bias also. So, even if we accept that near-future bias is irrational, that might be irrelevant to what our preferences for others ought to be. It might be that we still ought to care more about their present well-being than their distant-future well-being. If all this is right, the rest of my argument will follow—we will be morally required to be more concerned with other people's present interests than our own distant-future interests.


It might be objected that, if near-future bias is in fact irrational, it would be inappropriate to display other-regarding near-future bias. If I know as a matter of fact that someone's preferences are irrational, it seems that appropriate concern for them should mandate that I do *not* take on their irrational preferences. For instance, if my

¹¹ See Greene and Sullivan (2015), Sidgwick (1884), Rawls (1971), Lewis (1946), Nagel (1970), Broome (1991) and Brink (2011).

child has irrational preferences which I know would make their life worse overall, it seems that my appropriate concern for them as a parent should entail that I have preferences which contradict their irrational preferences, precisely because those preferences are irrational. So, if near-future bias is irrational, it could be argued that we should not display other-regarding near-future bias, but rather, remain temporally neutral with regards to other people's interests.

To respond, it could be said that such a mandate *not* to take on the irrational preferences of our children reflects particular aspects of the parent-child relationship. On the whole, it does not seem so strange to me that we ought to take on some of people's irrational preferences when acting on their sake. For example, a doctor might regard a patient's medical choices to be irrational, but they will still take on their irrational preferences out of respect for their bodily autonomy. In the same way, it could be argued that out of respect for individuals' opinions, we ought to take on the irrational preferences of other people and display other-regarding near-future bias.

However, even if it is the case that we ought to remain temporally neutral when it comes to other people's interests, it remains that we do display near-future bias regarding our own interests. While it is relatively easy to be temporally neutral with regards to the interests of distant strangers, we are naturally inclined to be near-future biased when it comes to our own well-being. If I am near-future biased regarding my own interests, but remain temporally neutral regarding other people's interests, the graph comparing the level of concern I have for myself and others over time would look like this:



As with the previous graph, we can see here that, given my near-future bias regarding my own interests, the distance between the level of concern I have for my distant-future interests and the interests of a stranger (distance b) is smaller than the distance between the level of concern I have for my present interests over the interests of a stranger (distance a). Again, this implies that, even if I am not required to be equally concerned for a distant stranger's well-being over my well-being at any given time, I morally ought to be willing to sacrifice some unit of my distant-future well-being in order to greatly increase a stranger's present well-being. This is because, if I am not willing, it implies that I am much too disinterested in the well-being of distant strangers, putting my own interests way above their interests than what is permitted by morality.

So, if we accept the irrationality of near-future bias, but we just happen to be near-future biased regarding our own interests, then my conclusion is more nuanced. We ought to either give up our near-future bias, or at least sometimes sacrifice our distant-future interests for other people's present interests. If it turns out, as I suspect, that near-future bias is psychologically very difficult to give up, then, in practice, most people will have to take the latter of these options. Given that we care less about our distant-future well-being, we should be willing to sacrifice a greater unit of our distant-future well-being than we would be required to sacrifice if it were our present well-being, in order to alleviate the present suffering of strangers.

5.2.2. *The Rationality of Near-Future Bias*

Another way to respond to the objection is to argue for the rationality of near-future bias by appealing to a psychological view about personal identity, such as the view defended by Parfit. Parfit argued that 'psychological connectedness' is one of the criteria which makes me the same person over time, by which he means the holding of direct psychological relations between myself at one point and myself at another. These psychological relations refer to psychological features such as memories, intentions and desires. Parfit claims that how much I care about my future self might depend on the strength of the psychological relations between me now and myself in

the future.¹² He argues that ‘since connectedness is nearly always weaker over long periods, I can rationally care less about my further future’. In other words, it can be rational to care less about myself in the distant future than myself in the nearer future, because time diminishes the degree of psychological connectedness. Similarly, Jeff McMahan argues that the ‘egoistic concern’ that it is reasonable for someone to have now for their own future good depends on the strength of the psychological connections holding between the person now and the person at the time at which the good will be experienced.¹³ Again, the strength of these psychological connections will be weaker over long periods of time, so it seems reasonable for someone to care less about their far-distant future.

Although appealing to Parfit or McMahan’s view may justify the rationality of near-future bias when it comes to our own interests, this kind of response does not justify other-regarding time bias.¹⁴ Appealing to psychological connections explains why we should care more about our *own* near-future selves, but why should we care more about other people’s near-future selves over their distant-future selves?

My answer to this challenge should now be familiar. If near-future bias is rational, our appropriate concern for other people mandates that we also take on their preferences. If my egoistic concern for myself dictates that I display self-regarding near-future bias, my appropriate concern for others should dictate that I display other-regarding near-future bias. It follows that, if it is rational for us to be near-future biased, and we are as concerned with the well-being of other people as we morally ought to be, then we morally ought to care more about other people’s present interests than about our distant-future interests.

5.2.3. *Pure Time Preference Versus Mere Time Preference*

Even if we reject the idea that psychological connectedness gives us grounds to be near-future biased, we can argue for the rationality of near-future bias by distinguishing

¹² Derek Parfit, *Reasons and Persons*, 1984, p. 313.

¹³ Jeff McMahan, *The Ethics of Killing: Problems at the Margins of Life*, 2002.

¹⁴ Indeed, Todd Karhu (2022) argues that an *advantage* of justifying future bias on the basis of the egoistic concern relation is that it explains away the asymmetry in the way we tend to display future bias when it comes to our own interests, but not when it comes to other people’s interests.

between *pure* positive time preference and *mere* positive time preference. When we say that near-future bias is irrational, this is to say that it is irrational to show a positive *pure* time preference, which is a time preference for utility or well-being under conditions of certainty. For instance, suppose you are offered a choice between two goods at an early and a late time respectively, and that you are certain that the late good will be just as valuable to you when you get it as the early good when you get it. If you still prefer the early good, then you are displaying pure time preference.

Although it may be irrational to display a positive pure time preference, it seems that we should not be temporally neutral regarding our well-being in the real world. For instance, few would dispute that it is rational to prefer near-term goods to far-term goods in real life. This is because we live in an uncertain world, where many things could happen which would make the far-term good less valuable to us. When we factor this uncertainty into our present decision, it may be rational to discount the value of the far-term good relative to the near-term good. In the same way that we discount the value of far-term goods, it seems that we should also discount the value of our distant-future interests on the grounds of uncertainty. Under conditions of uncertainty, it is *not* irrational to show a positive time preference regarding my well-being. The future, after all, is promised to no one; I may die tomorrow, or the world might end in fifty years.¹⁵ So, it seems that in practice we should discount our future in the way that I have argued, even if, under conditions of certainty, it would be rational to be temporally neutral.

The same argument applies when it comes to other people's interests. They too live in an uncertain world, and so it is rational for them to display near-future bias when it comes to their own well-being. Since we should show appropriate concern for others by taking on their rationally permissible preferences, we should display other-regarding near-future bias.

To sum up, even if we grant that near-future bias is irrational, my argument will still go through, so long as near-future bias is psychologically difficult to give up. And there is reason to think that near-future bias is not irrational, both because we can

¹⁵ See Dale Dorsey "A Near-Term Bias Reconsidered". *Philosophy & Phenomenological Research* 99/2 (2019): 461-477.

appeal to psychological connections to explain why we should care more about our near-future selves, and because we live in conditions of uncertainty.

5.3. Misgivings About the Economics

There are certain economic facts that might seem to undermine my argument. In this final subsection, I address some misgivings one might have about my argument in light of these facts.

5.3.1. *Compounding Interest and Returns on Investments*

So far, I have argued that, given our near-future bias, we ought to sacrifice some unit of our distant-future well-being in order to greatly increase a stranger's present well-being. However, compounding interest and investments make things more complicated. If we choose to save, the interest we gain from the sum of our savings will compound over time. If I choose to invest my money, I will gain returns on the investment over time.

For example, if you choose to save \$100 rather than donate that sum of money to charity, the interest you earn from saving will compound over the years so that the sum of money is much larger in the far-distant future. Furthermore, if you choose to invest this money rather than simply save, this could potentially *vastly* increase your future wealth. Given these economic facts, even if you discount your distant-future well-being, it may be that you ought to save or invest now rather than donate your money to help alleviate the current suffering of strangers.

To respond, I think that any additional gains from compounding interest or investment returns will be offset by the diminishing marginal utility of wealth, both with regards to your own wealth, and your wealth comparative to those in absolute poverty.

First, let's consider your own wealth. The diminishing marginal utility of wealth will mean that every dollar or pound will matter to you less than the last because each incremental increase in wealth will provide you with a smaller incremental increase in utility. So, if you are already wealthy, the interest added up over the years or the additional wealth gained by investing will not increase your well-

being very much. At most, it will just allow you to have an even cushier retirement. So, even if any money we save today will be worth a lot more in the future, given that the extra money in our distant future will provide us with only a small amount of utility *and* given that we care less about our distant future well-being, it remains that we should direct our money towards poverty alleviation.

Second, the problem of compounding interest can be offset by the diminishing marginal utility of your wealth *in comparison* to those in absolute poverty. As I explained above, it costs very little to do a lot of good for distant strangers in absolute poverty. In addition, there are significant returns on current benefits experienced by those in poverty. For example, health interventions not only spare people from disease, but arguably have beneficial effects on human capital formation and economic development.¹⁶ So, even if investing or saving could greatly increase your fortunes in the distant future and do you some good, the impact of giving what you have now to those in poverty will be far greater.¹⁷ My argument shows that we should at least be willing to sacrifice some unit of my distant-future well-being in order to greatly increase a stranger's present well-being. So, again, even if any money we save today will be a larger sum of money in the future, we ought to direct it to famine relief instead, in order to greatly increase a stranger's present well-being.

5.3.2. *Optimal Saving Policy*

¹⁶ For example, productivity losses due to malaria are estimated as 2.6% of annual household income for Malawi (Ettling et al. 1994, 'Economic Impact of Malaria in Malawian Households'), 2-6% of GDP for Kenya, and 1-5% for Nigeria (Leighton & Foster 1993, 'Economic Impacts of Malaria in Kenya and Nigeria'). Also, "one study on malaria eradication in the U.S. South (which had malaria until 1951) and several countries in Latin America suggests that a child who grew up malaria-free earns 50 percent more per year, for his entire adult life, compared to a child who got the disease. Qualitatively similar results were found in India, Paraguay, and Sri Lanka." (Bhanerjee and Duflo, *Poor Economics*, p.45)

¹⁷ This may imply that we should reject Patient Philanthropy, the view that individuals should invest and later donate financial resources, instead of donating now. However, if the returns of investing are sufficiently great, and we could do a lot more good by choosing to invest now and donate later, then we could justify doing so.

Suppose I am considering giving \$100 to charity this month. I have argued that if I discount my future welfare, I can donate the \$100 and reduce my savings (rather than my present consumption) by \$100, so that the donation comes at the expense of my future self, whose welfare matters less to me than the present welfare of strangers. However, if I am already following an optimal saving policy in light of my discount rate, the discounted future utility I get from \$100 in additional savings will be roughly the same as the present utility I get from \$100 in additional present consumption. In other words, taking the \$100 out of my savings shifts the welfare cost to the future, which I care less about, but also makes the welfare cost larger in undiscounted terms, to a roughly offsetting degree. So, the cost of taking the \$100 out of my savings will be just as great as taking the \$100 out of my present consumption, even after accounting for the fact that I discount my future welfare.¹⁸

To respond, even though we do discount our future welfare, I do not think the saving policy that many people adopt reflects this near-future bias. People are not, in fact, rational economic agents, and do not follow an optimal saving policy. Many people, for instance, save or invest *despite their near-future bias*, due to social expectations and pressures. If this is right, then taking \$100 out of our savings is not going to be as costly as taking it out of our present consumption, because we are *not* following an optimal saving policy in light of our discount rate. Rather, we are saving more than we should, given that we *do* discount our future welfare and given that we rationally *ought* to discount our future welfare given our uncertainty about the future.

Failing to follow an optimal saving policy in light of our near-future bias would not be morally problematic if it were only our own welfare that we needed to consider. Saving up more than we rationally ought to would not be immoral—we would merely be preventing our present selves from enjoying additional welfare and giving more to our future selves instead. However, it is not only our own well-being we must consider, but the well-being of other people too, including distant strangers. When we save more than we rationally ought to in light of our discount rate, we are depriving not only our present selves, but also distant strangers. This *is* morally problematic, for it implies that we care too little about the present condition of distant strangers.

¹⁸ I thank Christian Tarsney for pressing me on this point.

6. Conclusion

In this paper, I argued that our near-future bias implies that we morally ought to be less concerned about our own distant-future well-being compared with the present well-being of other people. At the very least, the difference between our level of concern for our distant-future interests and the present interests of others should be small, so that we are morally required to sacrifice some unit of our distant future well-being in order to greatly increase a stranger's present well-being. This is the case even if we are permitted a certain degree of partiality towards our own interests and the interests of our loved ones. I argued that this observation is particularly relevant when it comes to our moral obligations to give to charity, as for most people of a certain standard of living, the decision to give to charity will not usually affect their present well-being, but rather affect their distant-future well-being. If I am right, this argument calls into question the morality of saving up to secure our distant future at the expense of the current suffering of those in poverty.

References

- Banerjee, A. V., and Duflo, E. (2021) *Poor Economics: A Radical Rethinking of the Way to Fight Global Poverty*. New York, PublicAffairs.
- Brink, D.O. (2011) Prospects for Temporal Neutrality. In *The Oxford Handbook of Philosophy of Time*, ed. Craig Callender. Oxford: Oxford University Press.
- Broome, J. (1991) *Weighing Goods: Equality, Uncertainty and Time*. Basil Blackwell.
- Caruso, Eugene, Daniel Gilbert, and Timothy Wilson. "A Wrinkle in Time: Asymmetric Valuation of Past and Future Events" *Philosophical Science* 19/8 (2008): 796-801.
- Dorsey, Dale. "A Near-Term Bias Reconsidered" *Philosophy & Phenomenological Research* 99/2 (2019): 461-477.
- Dorsey, Dale. "Prudence and Past Selves" *Philosophical Studies* 175/8 (2018): 1901-25.

- Dougherty, Tom. "Future-Bias and Practical Reason" *Philosopher's Imprint* 15/30 (2015): 1-16.
- Ettling M, McFarland DA, Schultz LJ, Chitsulo L. "Economic impact of malarian in Malawian households." *Trop Med Parasitol* 41 (1994): 74-79.
- Greene, Preston, Andrew Latham, Kristie Miller, and James Norton. "Hedonic and Non-Hedonic Bias Toward the Future" *Australasian Journal of Philosophy* 99/1 (2021a): 148-63.
- Greene, Preston, Alex Holcombe, Andrew J. Latham, Kristie Miller, and James Norton. "The Rationality of Near Bias toward both Future and Past Events" *Review of Philosophy and Psychology* (2021): 905-922.
- Greene, Preston and Meghan Sullivan. "Against Time Bias" *Ethics* 125 (2014): 947-970.
- Hare, Caspar. "Self-Bias, Time-Bias, and the Metaphysics of the Self and Time" *Journal of Philosophy* 104 (7): 350 -373.
- Hare, Caspar. "A Puzzle About Other-Directed Time Bias" *Australasian Journal of Philosophy* 86 (2008): 269-277.
- Hare, Caspar. (2013) "Time – The Emotional Asymmetry" *A Companion to the Philosophy of Time*, eds. A. Bardon and H. Dyke, Malden, MA: Wiley-Blackwell: 507-20.
- Karhu, Todd. "What Justifies Our Bias Toward the Future?" *Australasian Journal of Philosophy*.
- Leighton C and Foster R. (1993) "Economic Impact of Malaria in Kenya and Nigeria," Heal Financing and Sustainability Project 6, Bethesda, Abt Associates, Maryland.
- McMahan, Jeff. *The Ethics of Killing: Problems at the Margins of Life* (2002). New York, NY: Oxford University Press.
- Nagel, T. (1970) *The Possibility of Altruism*. Oxford: Clarendon.
- Parfit, Derek. "Personal Identity" *The Philosophical Review* 80, no 1 (1971): 3-27.
- Parfit, Derek. *Reasons and Persons*. (1984) Oxford: Clarendon Press.
- Rawls, John. (1971) *A Theory of Justice*. Belknap: Cambridge, MA.
- Singer, Peter. "Famine, Affluence, and Morality" *Philosophy & Public Affairs* 1, no. 3 (1972): 229-243.