

Rationality with respect to people, places, and times

Larry S. Temkin 

Department of Philosophy, Rutgers University, New Brunswick, NJ, USA

ABSTRACT

There is a rich tradition within game theory, decision theory, economics, and philosophy correlating practical rationality with impartiality, and spatial and temporal neutrality. I argue that in some cases we should give priority to people over both times and places, and to times over places. I also show how three plausible dominance principles regarding people, places, and times conflict, so that we cannot accept all three. However, I argue that there are some cases where we should give priority to times over people, suggesting that there is impersonal value to the distribution of high quality life over different times.

ARTICLE HISTORY Received 23 October 2015; Accepted 13 November 2015

KEYWORDS Practical rationality; impersonal value; decision theory; temporal neutrality; impartiality; dominance principles

1. Introduction

Many believe that, regarding rationality, there are important respects in which we should treat people, places, and times the same. I think this is correct. On the other hand, many of the same people also believe that, regarding rationality, there are important respects in which we should treat people, places, and times differently. I think this, too, is correct. This raises the obvious question of under what circumstances, and in what respects, we should, or should not, treat people, places, and times the same. This is, I believe, an extremely important, but underexplored, question which raises a host of rich, complex, and thorny issues. In this article, I will begin the difficult task of addressing this question.

The aim of this article is not to *settle* the question of how we should think about people, places, and times for the purposes of practical reasoning. Far from it. Such a task lies well beyond the scope of a single article. Instead, I wish to explore a number of issues pertinent to this topic. I cannot emphasize enough the *preliminary* nature of my exploration. I am acutely aware that many of my

CONTACT Larry S. Temkin  ltemkin@rci.rutgers.edu

© 2015 Canadian Journal of Philosophy

arguments and examples are merely suggestive, rather than conclusive; that not everyone will share my intuitions about the examples I give; that some who share my intuitions will offer deflationary accounts as to why such intuitions are not to be trusted; that some will draw different conclusions than I do regarding what the examples and our intuitions about them suggest; and that some of my arguments and examples are open to serious worries or objections that I cannot adequately deal with here. In addition, my article makes no attempt to be complete, or to offer an even-handed treatment of the issues.

In Section 2, I'll offer just a few of the many considerations that might be offered in support of the view that in certain respects, we should treat persons, places, and times the same in our practical deliberations. In Sections 3–5, the bulk of the article, I will offer considerations – some of them obscure and many of them controversial – in support of the view that often we may, and sometimes must, treat people, places, and times differently in our practical deliberations. In Section 6, I'll note a number of related issues that still need to be addressed. Despite all these hedges and qualifications, I hope to convince the reader that the issues raised in this article are interesting and important, and that they require much more attention than they have heretofore been given.

One final caveat. Many of the examples presented in this article involve moral considerations, many of which are axiological in nature. But I hope it is clear that they have implications for the domain of non-moral practical reasoning, as well as the domain of practical reasoning for which moral considerations are pertinent.

2. Stage setting

The view that there are important respects in which we should treat people, places, and times the same for the purposes of practical reasoning is supported by an ingenious argument of Parfit's (1984, Part Two). Parfit contrasted three different positions, each with a different structure: a *present-aim theory*, which someone like Williams (1981a, 1981b, 1985) might adopt; a *self-interest theory*, which has been a dominant conception of individual rationality since the early Greeks (Plato, *The Republic*); and a *neutralist theory*, favored by classical utilitarians like Mill (*Utilitarianism*) and Sidgwick (1907). Parfit offered an intriguing strategic metaphor, suggesting that the self-interest theory occupied an indefensible 'no man's land' between the other two positions. Parfit noted that the present-aim theory is a 'pure' theory, *relativizing* the reasons one has to a particular person, place, and time.¹ On such a view, for the purposes of practical deliberation, only an agent's *own* interests matter, and only those that the agent has *here* and *now*, at the time and place where the practical deliberation is occurring. Likewise, neutralist theories are 'pure,' in that, for the purposes of practical deliberation, they treat *all* people, places, and times *neutrally*, giving equal weight to the interests of all people at all places and all times.² By contrast,

the self-interest theory is a 'hybrid' theory, it *relativizes* with respect to *people*, telling each person that she only needs to give weight to *her* interests, but it is *neutral* with respect to *space* and *time*, holding that a person should give *equal* weight to *each* moment of her life, no matter where, or when, it occurs.

Parfit suggested that the self-interest theory's hybrid position may be unstable and indefensible. In particular, he claimed that the strongest arguments that the self-interest theory might muster in opposition to the present-aim theory, against relativizing with respect to place and time, might analogously tell against relativizing with respect to individuals, and so may carry one *from* the present-aim theory, *beyond* the self-interest theory, and all the way *to* the neutralist position. Likewise, Parfit suggested, the strongest arguments that the self-interest theory might muster in opposition to the neutralist position, against being neutral with respect to *people*, might analogously tell against being neutral with respect to place and time, and so may carry one *from* the neutralist position, *beyond* the self-interest theory, and all the way *to* the present-aim theory.

Thus, Parfit suggested, there are powerful reasons for any theory of practical reasoning to be 'pure,' and to treat people, places, and times alike: either we should be *relative* with respect to all three, or we should be *neutral* with respect to all three, but what we should *not* do is to treat people, places, and times *differently* insofar as we believe that reasons are either relative or neutral.

Many theorists accept the view that we should treat people, places, and times the same, and that in fact we should be *neutral* between them. This is the view of many consequentialists, including all of the classical utilitarians.

The view that rationality requires us to be *neutral* with respect to time was expressed nicely by Henry Sidgwick, when he wrote that:

'Hereafter *as such* is to be regarded neither less nor more than Now.' ... the mere difference of priority and posteriority in time is not a reasonable ground for having more regard ... [for] one moment ... [over] that of another ... 'a smaller present good is not to preferred to a greater future good' (allowing for differences of certainty). (1907, Book III, 381)

Similar claims might be made with respect to people and places. Thus, echoing Sidgwick, one might hold, regarding persons, that:

Me *as such* is to be regarded neither less nor more than You. The mere difference of who is who – that I am I, and you are you – is not *itself* a reasonable ground for having more regard for one of us over the other. Hence, a smaller good for one person is not to be preferred to a greater for another (merely in virtue of the fact that each person is the person that he or she is).

Likewise, one might hold, regarding space, that:

Here *as such* is to be regarded neither less nor more than There. The mere difference of nearness or distance in space is not a reasonable ground for having more regard for one location over that of another. Hence, a smaller nearer good is not to be preferred to a greater further good.

Now, as an egalitarian, who also believes in certain agent-relative duties and permissions, I have never been attracted to the sort of 'pure' neutralist position of classical utilitarianism, whose sole focus is on *how much* utility obtains, without regard to how it is *produced* or *distributed* by, and across, different people, places, and times. Still, there is a powerful attraction to the kind of reasoning expressed by Parfit and Sidgwick, in support of the views that in *some* important respects, we should treat people, places, and times the same, and be *neutral* with respect to all three. Unfortunately, here, as elsewhere, the devil is in the details and, sadly, an adequate account of the details has not yet been given. More to the point, whatever kernel of truth there may be to the sort of views expressed by Parfit and Sidgwick, I believe that there is good reason to reject any blanket suggestion that we should treat people, places, and times the same, or be neutral between them, beyond the standard agent-relative objections that have been mustered against such positions.³ While I cannot fully defend my view here, in the remainder of this article, I shall present a few of the considerations that underlie my thinking about this matter. As we will see, the question of whether we should treat people, places, and times the same goes well beyond the issue of whether reasons should be neutral with respect to each, relative with respect to each, or neutral with respect to some but relative with respect to others.

3. Some musings about space and time, and worries about treating them the same

It may seem *obvious* that we should treat space and time the same. However, I'm not so sure about this, especially if it is supposed to be an a priori *truth* that holds regardless of the metaphysics of space and time. Suppose, for example, that we lived in a universe that extended infinitely in all directions, spatially, and infinitely towards the past and future, temporally. Suppose, further, that time's passage is not an illusion, and that time is 'directional,' such that the past is receding at a constant rate from the steadily changing present, even as the future is steadily moving at the same constant rate towards the present.⁴ *Perhaps* such a view is incoherent, or metaphysically impossible, but if not, must we treat space and time the same for the purposes of practical deliberations?

Consider the following thought experiment. Suppose I learn that our civilization will live in our galaxy another 1000 years, and then die out. I also learn that in a distant galaxy, another civilization will exist for the same 1000 years, and then die out. I then learn that this is also the case in some third and fourth distant galaxies. I find this all quite interesting. It is somewhat *pleasing* to me to learn that there are, in fact, advanced civilizations living in galaxies far away.

Next, suppose I also learn that beyond the fourth galaxy, there is nothing but cold, empty, space. This, too, I find interesting. However, I must confess that learning that fact bothers me some, but not very much. Indeed, if events beyond the fourth galaxy were about to unfold which would make those distant reaches

inhospitable to all life forms in perpetuity, I wouldn't think it especially important for us to make significant sacrifices, if we could, to prevent that from happening.

Suppose, on the other hand, I vary the story. As before, I learn that civilization in our galaxy will die out in 1000 years, but I learn that after ours dies out, another civilization will arise and persist for 1000 years in a second galaxy. I also learn that this will happen again, a third and fourth time. But after that, I learn, there will be *nothing* but cold, empty, space, *forever*. For some reason, *that* knowledge would bother me *a lot*. Indeed, if events were about to unfold which would make the universe uninhabitable for any life forms 4000 years from now, unless we made significant sacrifices to prevent that from happening, I would feel quite *strongly* that we should do so, and I would feel that way even if I knew that *our* civilization was going to die out in 1000 years, and that the distant future civilizations would do nothing to further *our* particular hopes, projects, or ideals.⁵

Here is a variation of the example. Suppose that I am living in a world where the only sentient beings are human, where our civilization will have persisted for a total of 10,000 years before dying out, and where there are 10 billion people alive during each period where our civilization exists, each of whom is at a high level, *h*. I believe that no other civilizations exist elsewhere in space, and that no other civilizations existed before ours, or will exist after ours. I then learn that I am mistaken in one of two ways. Either I am mistaken about there being no other civilizations in space, and in fact there are 10,000 other planets that will be populated by 10 billion beings also at level *h* during the same time period that our planet is populated, but I am right that no other civilizations will exist in the universe prior to, or after, our civilization; or, alternatively, I am mistaken that no other civilizations exist during another time period, and in fact that are 10,000 other civilizations of 10 billion beings, each of whom is living at level *h* in a distinct non-overlapping 10,000 year time period of its own, but I am right that there are no other civilizations living elsewhere in space during the time period where our civilization persists. Here, I believe it would be much better if I were mistaken in the second way than the first. And that is because I believe it would be much better for there to be 100,010,000 years where different groups of 10 billion sentient beings are living at a high level, than for there to be only 10,000 years where the universe is occupied by sentient beings with high-level lives, even if, during those 10,000 years, there would not merely be 10 billion people alive, but 100,010 billion people alive.

Note, in both cases, there would be the same sum total of utility. Indeed, in both cases, the *very same* people might exist at the *very same* levels of well-being. Still, I believe that the alternative where the many people with high levels of well-being are dispersed throughout time, so that there are lots of cases of high-quality lives stretched over many eons, is better than the alternative where the many people with high levels of well-being are dispersed throughout space, so that there are lots of cases of high-quality lives stretched over many miles, or acres.

Let us add another variation to the example. Suppose that I was mistaken in both respects. The reality is that our civilization is at the end of a long line of 10,001 contemporaneously populated planets, like ours, extending deep into space along a single ray from the center of our planet, and, in addition, it is the last of 10,001 non-overlapping civilizations, also like ours, extending deep into the past for roughly 100,000,000 years. I then learn that long ago, an infallible predictor set matters in motion that would determine how the universe would unfold, depending on the actions of the members of our planet.

If we took certain steps, A, which would lower the quality of our lives by 20%, then, in fact, there would be an additional 10,000 contemporaneous civilizations extending deep into space along another ray from the center of our planet, each of whose members would be at level *h*. If we didn't do A, then there would be no other planets elsewhere in space populated by sentient beings during our civilization's existence, other than the 10,000 contemporaneously populated planets about which I already know. In addition, I learn that if we took certain other steps, B, that would also lower the quality of our lives by 20%, but, in that case, there would be an additional 10,000 non-overlapping civilizations, like ours, each of whose members would be at level *h*, extending forward in time for a total of another 100,000,000 years after our civilization dies out. If we don't do B, the universe will remain utterly devoid of all sentient life once our civilization, and the other 10,000 contemporaneous civilizations, come to an end.

Faced with the knowledge of these alternatives, I think it would be important that our civilization took steps B, to ensure that high-level sentient life persisted in the universe for another 100,000,000 years, and that it would almost certainly be wrong of us not to do so. I think it would be much less important for our civilization to take steps A, to ensure that there be even *more* high-level civilizations living in space at the same time as ours, *in addition* to the 10,000 other such civilizations that will *already* be existing elsewhere in space during that time. Moreover, I think it would probably not be wrong of our civilization to fail to do A.

My own view about this case is that the universe might well go *best* if we did *both* A and B, go second best if we did B but not A, go third best if we did A but not B, and would go worst if we failed to do either A or B. Hence, I am *not* denying that, other things equal, it might be important to populate different regions of space with high-quality lives. But the key point, for my present purposes, is that I think we should treat time and space *differently*, in this context. Specifically, I believe that, in certain cases at least, we should give *greater* priority to filling differing periods of *time* with quality life, than to filling different locations in *space* with quality life.⁶

Is it crucial to this example that the different periods of time to be filled come *after* our civilization will die out? Not to my mind. I would feel similarly about the greater importance of doing B, rather than A, if as a result of our doing B, which would lower the quality of our lives by 20%, there would be an additional

10,000 non-overlapping civilizations, like ours, each of whose members would be at level h , extending *backward* in time for a total of *another* 100,000,000 years *before* the 100,000,000 years of civilization that preceded our civilization's existence. Thus, my thought is not merely that it is more important for high-quality life to be dispersed into the *future*, rather than to be dispersed across *space*, but rather that, more generally, it is more important for high-quality life to be dispersed across *time* rather than across *space*.⁷

Here is another example. As before, suppose, that our civilization consisted of 10 billion people, all at the high level of h , and that altogether our civilization would last 10,000 years, before dying out. I then learn that we were the first sentient beings to exist in the universe, and that no other sentient beings will exist for another 7,000,000,000,000,000,000,000,000 (7 octillion) years, after which there will be one more sentient civilization, like ours, with 10 billion people that will last for 10,000 years, before the universe grows cold, forever. I would, of course, feel incredibly lucky to know that we were one of only two sentient civilizations that would ever exist in the whole duration of the universe. But I would also feel that the overall goodness of the universe was a *pittance* compared to what it might have been. I would think it a great cosmic tragedy that the vast majority of time periods in the universe were utterly devoid of high-quality sentient life. Correspondingly, if our civilization could somehow find a way to ensure that high-quality sentient life would continue for the seven *octillion* years after our civilization dies out, I believe there would be *powerful* reason for us to do so, and I believe this even if it would require *substantial* sacrifice on the part of our civilization to bring about the better outcome.

Consider, next, the following. If one travels by plane from one end of the US to the other, and gazes out one's window, one may be struck by the fact that outside of a few major metropolitan areas, much of the US consists of vast unpopulated tracts of land. The same is true for Canada, Australia, Russia, and much of Africa, Asia, South America, and Northern Europe. Iceland is almost devoid of people; Greenland, the Arctic, and Antarctica even more so. Moreover, 71% of the Earth's surface is water. When I think about these facts, I *don't* think that it is a great cosmic tragedy that the vast majority of spatial locations, right here on Earth, are utterly devoid of high-quality sentient life. Should I?

Nor do I believe that my reaction here is solely due to the assumption that if all those spaces were filled with sentient life, none of it would be of high quality. Even assuming that *everyone* who would live on the Earth would have a high-quality life, I don't see a compelling reason to increase the size of the Earth's human population from 7 billion to, say, 100 billion or more, even if we could.

Or let us turn our gaze inward, for a moment. It is estimated that 1% of every atom is composed of protons, neutrons, and electrons, and that the other 99% is empty space.⁸ And it is estimated that the average human adult has 7,000,000,000,000,000,000,000,000 (7 octillion) atoms in his or her body. There

are over 7 billion humans on the Earth. Given all this, should we regard it as a great cosmic waste that the *vast majority* of spatial locations within each human body is devoid of high-quality sentient life? Would our world, or the universe, be much better if within each atom of each human being (not to mention all the other atoms on Earth!), there were subatomic sentient beings possessing high-quality lives? If, contrary to fact, our civilization could somehow find a way to create subatomic sentient beings with a high quality of life, would there be powerful reason for us to fill each of the 7 octillion atoms of each 'average' adult human, as well as all of the other innumerable atoms of the rest of the 7 billion members of the human population, as long as when doing so our own quality of life remained the same? Would there be powerful reason to bring about such an outcome, even if doing so required a *substantial* sacrifice on the part of our civilization?

Suffice it to say, when I think about *all* the locations of space on Earth devoid of sentient beings with high-quality lives, I'm not moved in anything like the way I am when I think of vast regions of time that are devoid of such beings. In sum, in general, I think it important that many *times* be filled with flourishing beings, but not nearly as important that many *spaces* be filled with flourishing beings.⁹ Perhaps I'm mistaken about all this, of course. But, for now, at least, I see no compelling reason to abandon my views about this matter.

I have been focusing on cases involving the high-quality lives of sentient beings. Unsurprisingly, my judgment would flip regarding the relevant importance of filling time vs. space if the lives in question were miserable – well below the level at which life ceases to be worth living.

Suppose, for example, that there are two ways in which the universe might unfold. In one, there are 10 billion planets, each populated with 10 billion people, each of whom lives for 100 years, and all of whom are suffering unrelenting agony. In the other, there are 10 billion planets each populated with 10 billion people, each of whom lives for 100 years, and all of whom are suffering unrelenting agony. Assume that the very same people exist in each universe, and that no one else exists, other than the 100,000,000,000,000,000,000 people in sheer agony. From the *subjective* standpoint of each individual, it won't matter which universe exists. And from the standpoint of total disutility, each universe will be identical. Suppose, however, that in the first alternative, each person lives *simultaneously*, while in the second alternative, each planet is populated during a different time period. In that case, I think the second alternative is *worse* than the first. I think a universe where there is *vast* suffering, but where the suffering only lasts for 100 years, is much better than a universe where there is the same *total* amount of vast suffering, yet there are great numbers of people in great agony not merely for 100 years, but for 1000 trillion years. So, in my judgment, it is bad if many *spaces* are filled with agonizing lives that are worth *not* living, but it is *worse* if many *times* are filled with agonizing lives that are worth *not* living. Thus, the relative importance of filling times, or spaces, with sentient

beings, depends on whether the value of those being's lives are worth living, or worth *not* living.¹⁰

Let me turn, next, to a different point. Suppose that God were deciding to populate an infinite number of planets, and time periods of 100 years each, with sentient beings whose lives were *miserable*. Her plan is to have 10 billion miserable beings living on planet one billion and one, during time period one billion and one, 10 billion *different* miserable beings living on planet one billion and two, during time period one billion and two, 10 billion *different* miserable beings living on planet one billion and three, during time period one billion and three, and so on, for all of eternity. Just before doing so, God decides that She will create the *very same* miserable beings as She was originally intending to, and that each of them will live during the *very same* time periods as She was originally intending for them, but that She will *shift* which particular *planets* they occupy, so that *each person's place* in space would be different.

Specifically, suppose that God decides to put the people who *would* have occupied planet one billion and one on planet one, instead, the people who *would* have occupied planet one billion and two on planet two, instead, the people who *would* have occupied planet one billion and three on planet three, instead, and so on. One might, if one likes, imagine that each of an infinite number of planets are spaced an equal distance apart, say k miles, along an infinite, straight, Euclidean line, so that God's choice involves placing each person on the planet She originally intended for them, or, instead, shifting each person k billion miles in the same direction along the line of planets to a different planet. To my mind, the difference between these two prospects has *no* moral significance. As between *these* options, where everything is the same except for *where* in *space* the infinite people lived, I would be *utterly* indifferent.

Suppose next, however, that God decides that She will create the *very same* miserable beings and place them on the *very same* planets as She was originally intending to, but that She will shift *when* they live. Specifically, suppose that God decides to put the people who *would* have occupied time period one billion and one in time period one, instead – where, we are assuming, time period one begins one billion years *earlier* than time period one billion and one – that She decides to put the people who would have occupied time period one billion and two in time period two, instead, the people who would have occupied time period one billion and three in time period three, instead, and so on. I find the difference between *these* two prospects to be morally significant. Notwithstanding the so-called 'fact' of infinity that tell us that, over the course of time, there will be just *as much* miserable existence in each of the two alternatives, I believe there is reason to favor God's original plan, over Her revised plan, in which miserable existence will begin one *billion years earlier*, and then continue, unabated, afterwards.

If you aren't convinced by the previous example, consider a variation that would directly affect you. Suppose that God tells you that She is planning to send

you to Hell, where you will suffer unbearable pain. She then tells you that once you are sent to Hell you will remain there forever, but that She is willing to give you some choice as to *when* your agony begins. You can either start immediately, or you can start sometime later; however, if you choose to start later, God will put you in a state of suspended animation between now and when you start, so that you will experience nothing between now and when you begin your torments in Hell. Dismayed, you begin asking how long you can delay it. A week? Sure. A Month? No problem. What about a year? A decade? A century? God is happy to go along with any of those choices. You decide to be bolder. You ask if you can delay by a *million* years. She agrees that that, too, is possible. You decide to be bolder still. Having recently learned how big an octillion is, you ask God if you could delay your entry to Hell by an octillion number of years. At this point, you have tried even God's (infinite!) patience, and She replies that yes, indeed, you can delay your entry by that much, but no longer!

At this point, what would you decide? Would you reason that since you won't be gaining anything positive by delaying, and will be spending an eternity in Hell once you're there, it doesn't matter *when* you start, since at the end of time, as it were, you will have spent *just as much* total time in Hell? Or would you choose to enter a state of suspended animation, and delay your entry as long as possible, taking the option of entering Hell in an octillion years? I know what I would choose, for myself, or anyone else that I dearly loved. I would choose the *latest possible* entry date that God permitted, and I believe that it would be perfectly rational for me to do so, and *irrational* for me *not* to do so!¹¹

Suppose, next, that God tells you that she is planning to send you to Hell immediately. She further informs you that Hell consists of an infinite number of planets spaced an equal distance apart, say k miles, along an infinite, straight, Euclidean line, each of which is labeled by an integer. As it happens, she is planning to place you on planet 1, where you will remain for one year, after which you will be on planet 2 for a year, then planet 3 for a year, and so on, for all of eternity. However, if you want, you have the option of entering Hell on any of the other planets located along the infinite straight line, where you will remain for one year, after which you will be moved to the next highest numbered planet where you will again remain for a year, and this pattern will continue for all of eternity. You immediately ask if any of the planets are less torturous than the others. She assures you that they are not. They are all equally torturous. At that point, would you bother to shift your entry point into Hell to a different planet at a different location in space? Would you try to bargain with God to please let you move your entry point from planet 1 to planet 2, planet 100, planet 1 million, or planet 1 octillion? I see no reason why one would. If, in fact, Hell's planets are all equally bad, it seems clear that there would be *no* rational basis for preferring to be in one particular location in space rather than any other.

I realize that these past two cases involve infinity, and that our intuitions about such cases are notoriously problematic. Nevertheless, I don't think my

views about these cases are implausible, and they suggest a further asymmetry about space and time for the purposes of practical reasoning. In some infinite cases, at least, merely shifting the *spatial* locations of sentient beings will be morally irrelevant, whereas shifting the *temporal* locations of sentient beings may be significant.¹²

4. Dominance principles with respect to people, places, and time

Parfit’s argument against the Self-Interest Theory suggests that rationality requires that we treat persons, places, and times the same in certain key respects, and, in particular, that if we should be neutral with respect to one, we should also be neutral with respect to the others. Reasoning along similar lines, it might seem natural to assume that *if* we accept a dominance principle with respect to *one* of these categories, then we should *also* accept a similar dominance principle with respect to the other categories. Consider, for example, the following three dominance principles regarding utility.

Spatial Dominance Principle: for any two alternative outcomes, A and B, if A and B involve the same regions of space, and A is better than B regarding utility in *every* region of space, then A is better than B regarding utility.

Temporal Dominance Principle: for any two alternative outcomes, A and B, if A and B involve the same regions of time, and A is better than B regarding utility in *every* time period, then A is better than B regarding utility.

Personal Dominance Principle: for any two alternative outcomes, A and B, if A and B involve the same people, and A is better than B regarding utility for *every* person who will ever live, then A is better than B regarding utility.

Intuitively, many would find each of the preceding dominance principles plausible. Moreover, as indicated, influenced by reasoning of the sort appealed to by Parfit, many might assume that if one of the dominance principles is true, then the others must also be true. But this assumption is clearly false. To see this, consider Diagram 1.¹³

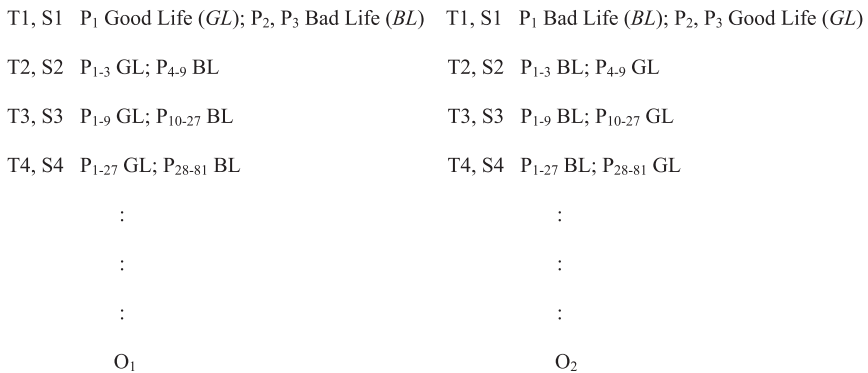


Diagram 1.

O_1 is one possible outcome. In that outcome, there is one person, P_1 , living in time period one, and spatial region one, who has a *good* life, well *above* the level at which life ceases to be worth living, but there are *twice* as many other people, P_2 and P_3 , who have *bad* lives, well *below* the level at which life ceases to be worth living. In time period two, P_1 through P_3 have moved to spatial region two, where they all enjoy good lives, but unfortunately, in that time period, and at that location, twice as many other people, P_4 through P_9 , have come into existence, and their lives are as bad as P_2 and P_3 's lives were during T_1 . In time period three, P_1 through P_9 have all moved to spatial region three, where they all enjoy good lives, but unfortunately in that time and location, twice as many other people, P_{10} through P_{27} , have come into existence, and their lives are as bad as P_2 and P_3 's lives were during T_1 . Outcome One continues to unfold, in this ever-expanding manner, forever, with each time period lasting for one day, and each person living for 100 years total, before dying. Here, and below, we assume that the positive value of each good moment is the same, the negative value of each bad moment is the same, and that the two values sum to zero, so that a life containing an equal number of moments of good and bad life will have a net value of zero, a life containing more moments of good life than bad life will have a positive net value, and a life containing more moments of bad life than good life will have a negative net value.

Outcome Two contains the *very same people* as Outcome One, P_1 , P_2 , P_3 , etc., and is analogous to, though the reverse of, Outcome One. Specifically, in Outcome Two, there is one person, P_1 , living in time period one, and spatial region one, who has a *bad* life, well *below* the level at which life ceases to be worth living, but there are *twice* as many other people, P_2 and P_3 , who have *good* lives, well *above* the level at which life ceases to be worth living. In time period two, P_1 through P_3 have moved to spatial region two, where they all suffer bad lives, but fortunately, in that time period, and at that location, twice as many other people, P_4 through P_9 , have come into existence, and their lives are as good as P_2 and P_3 's lives were during T_1 . In time period three, P_1 through P_9 have all moved to spatial region three, where they all suffer bad lives, but, once again, fortunately in that time and location, twice as many other people, P_{10} through P_{27} , have come into existence, and their lives are as good as P_2 and P_3 's lives were during T_1 . As before, Outcome Two continues to unfold, in this ever-expanding manner, forever, with each time period lasting for one day, and each person living for 100 years total, before dying.

Given our assumption that each moment of good life would balance equally against a moment of bad life, and the further moral assumption that in this example, there is no reason to favor one person over that of any other, how do Outcomes One and Two compare regarding utility?

On an *Impersonal Neutralist View*, of the sort favored by classical utilitarians, O_1 and O_2 might be judged as *equally good*. After all, if one is indifferent to *where* in space, time, or lives goods or bads are located, and *only* pays attention

to *how much total* good and bad exists in the world, then one may judge that O_1 and O_2 are equally good, since, in each outcome, there would be an infinite number of good days, and an infinite number of bad days lived, and in each case the orders of infinity of the number of good and bad days would be the same.

Alternatively, if one compares the two outcomes place by place, or moment by moment, Outcome Two would be clearly *better* than Outcome One, in accordance with the Spatial and Temporal Dominance Principles. This is because for *every* spatial region, S_n , and *every* temporal region, T_n , there will be twice as many people with good lives as with bad lives in Outcome Two, while there will be twice as many people with bad lives as with good lives in Outcome One.

So, should we conclude that, regarding utility, Outcomes One and Two are *equally good*, in accordance with the Impersonal Neutralist View, or that Outcome Two is *better than* Outcome One, in accordance with the dominance principles with respect to space and time? I find such judgments *very hard* to believe. After all, in O_1 , each person has exactly *one* bad day, and the *rest of his or her 100 year life* is good. In O_2 , on the other hand, each person has exactly *one* good day, and *the rest of his or her 100 year life* is bad. I know which of these outcomes I would want to obtain for myself, a loved one, or anyone else who was not pure evil!

O_1 is a world where *everyone* has lives that are very good *every single day but one*. O_2 is a world where *everyone* has lives that are very bad *every single day but one*. Clearly, *every* member of O_1 has a life which is, on balance, *well worth* living, whereas *every* member of O_2 has a life which is, on balance, *well worth not* living. Given all this, I firmly believe that Outcome One is *better* than Outcome Two, in accordance with the Personal Dominance Principle regarding utility.

In this example, we can accept the dominance principle regarding *people*, or we can accept the dominance principles regarding *space and time*, but we *cannot* do both! Here we have a *proof* that, unless we reject *all three* dominance principles, in some cases, at least, we *should* not, and *cannot*, treat space and time the same way as we treat people. So, *should* we reject all three dominance principles, in favor, perhaps, of the Impersonal Neutralist View? I don't see why. At least in the sort of case that we have been considering, the Personal Dominance Principle seems clearly true!^{14,15}

The preceding argument suggests that, for certain cases, at least, we should give priority to distributions of well-being across people over distributions of well-being across time. And earlier, I suggested being more concerned about distributions of well-being throughout time, than throughout space. The priority rankings of people over both time and space, and time over space, for some cases, at least, might be further buttressed by considering Diagram 2.

In Diagram 2, O_1 and O_2 are outcomes with an infinite number of people, P_i or Q_j , with each person, located at a particular location in space, S_k , and a particular location in time, T_l , at a level corresponding to one of the integers. So, for example, in Outcome One, person P_0 is at level 0, at temporal location 0

.... $T_{-4}, T_{-3}, T_{-2}, T_{-1}, T_0, T_1, T_2, T_3, T_4, \dots$
 $S_{-4}, S_{-3}, S_{-2}, S_{-1}, S_0, S_1, S_2, S_3, S_4, \dots$
 $P_{-4-4}, P_{-3-3}, P_{-2-2}, P_{-1-1}, P_{00}, P_{11}, P_{22}, P_{33}, P_{44}, \dots$

O_1

.... $T_{-3}, T_{-2}, T_{-1}, T_0, T_1, T_2, T_3, T_4, T_5, \dots$
 $S_{-5}, S_{-4}, S_{-3}, S_{-2}, S_{-1}, S_0, S_1, S_2, S_3, \dots$
 $Q_{-4-4}, Q_{-3-3}, Q_{-2-2}, Q_{-1-1}, Q_{00}, Q_{11}, Q_{22}, Q_{33}, Q_{44}, \dots$

O_2

Diagram 2.

and spatial location 0, while in Outcome Two, person Q_{-4} is at level -4 , at temporal location -3 and spatial location -5 . For the purposes of this example, I am assuming that the metaphysics of space and time allow for the identification of the same spatial and temporal locations across different possible outcomes, so that for each k and l , S_k corresponds to the very same location in space in each outcome, and T_l corresponds to the very same location in time in each outcome. If such an assumption is coherent, then Diagram 2 illustrates that, for some cases, at least, the Spatial and Temporal Dominance Principles are incompatible with each other. Thus, for such cases, we can reject both, but we can't accept both. This is because, in Diagram 2, Outcome One is *better* than Outcome Two at every point in *time*, but it is *worse* than Outcome Two at every point in *space*.

Now assume, temporarily, that the populations of the two outcomes are wholly distinct. In that case, I can see why someone might claim that each outcome is *equally* good, since each involves an infinite number of people, such that for each integer, there is exactly one person whose level of well-being is accurately represented by that integer. In that case, one would be rejecting *both* the Spatial and Temporal Dominance Principles. But my own judgment, in this case, is that we should *accept* the judgment yielded by the *Temporal* Dominance Principle, and *reject* the judgment yielded by the *Spatial* Dominance Principle. That is, in this case, I would judge Outcome One as better than Outcome Two, since it is better at each moment in time, and, to my mind, there is neither a compelling reason to ignore this consideration, nor a countervailing reason outweighing it.

But, of course, as the previous discussion makes plain, I believe that there *could* be such a reason. In particular, if the *same* people would exist in each outcome, and they would *each* be better off in one of the outcomes than the other, then, in accordance with the *Personal Dominance Principle*, I would regard the outcome in which they were all better off as better than the other outcome, regarding utility, regardless of how the two outcomes compared in accordance with either the *Spatial* or *Temporal Dominance Principles*.

5. Avoiding cases involving infinity

Some people will worry about some of my examples because they involve appeals to infinity about which our intuitions are notoriously unreliable. For those who have such worries, let me make several comments.

First, I believe that one must distinguish between different kinds of examples involving infinity. Some rely on moves that are clearly dubious, for example, when one 'reorders' the different members of an infinite sequence, say, by 'moving forward or backwards' certain members of the sequence, but not others, in order to shift our intuitions about the overall value of the infinite sequence, either when considered by itself, or in comparison with some other sequence. I fully *agree* that any intuitions that we might have about such 'reordered' infinite sequences are not to be trusted, but I note that *none* of my appeals to infinity involve such dubious moves. Indeed, I believe there is nothing 'tricky,' artificial, or dubious about the infinite sequences that I have considered in this article, that provides good reason to doubt our intuitions or judgments about them. To the contrary, although I cannot pursue this further, here, I believe that there are good reasons to accept our judgments about the various examples I have invoked involving infinity.

Second, I remind the reader that not all of my arguments involved appeals to our intuitions about infinite cases. Many of my arguments focused on finite cases.

Third, while I find some of my cases involving infinity particularly compelling – which is why I employ them – I believe that my main conclusions *could* have been argued for without appealing to such examples. In particular, I believe that there are a host of strong reasons, that don't appeal to infinity, to worry about any 'pure' neutralist position requiring us to be strictly neutral between any 'locations' of people, places, or times at which utility might obtain. I also believe that there are some finite cases where we should reject the rankings generated by the *Spatial* and *Temporal Dominance Principles*. In support of these claims, I offer the following observations.

First, consider the widely held view that the *shape* of a life matters, so that a life that begins poorly, but steadily improves, and ends well, would be better than a life that begins well, but steadily declines, and ends poorly, *even if* the

two lives contained the same *total* amount of well-being. Clearly, this view is not strictly neutral as to *when*, in a life, well-being occurs.

But notice, most people holding this view regard space and time differently in this respect. They *don't* believe, for example, that if someone was at level 100 for the first 20 years of his life, at level 400 for the next 20 years, at level 700 for the next 20 years, and at level 1000 for his final 20 years, that the *overall* quality of his life would vary depending on *where*, in space, he lived *during* those different periods. Such a life would be *equally* good if the person lived in *one* place all his life, or in *different* places. Indeed, as long as the levels for each period weren't affected, he could move to any combination of different *places*, *in any order*, without affecting the overall quality of his life. So, this is a non-infinite example where many believe that space and time should be treated differently, as the ordering of well-being in *time* seems *relevant* for the overall quality of a life, in a way that the ordering of well-being in *space* does not.

Next, consider three principles that I presented in my book, *Rethinking the Good*:

The Second Standard View – Trade-offs between Quality and Number Are Sometimes Undesirable Even When Vast Numbers Are at Stake: If the quality of one kind of benefit is 'sufficiently' low, and the quality of another kind of benefit is 'sufficiently' high, then an outcome in which a relatively small number of people received the higher quality benefit would be better than one in which virtually any number of people received the lower quality benefit. (Temkin 2012, 32)

The Disperse Additional Burdens View: in general, if additional burdens are dispersed among different people, it is better for a given total burden to be dispersed among a vastly larger number of people so that the additional burden any single person has to bear within her life is 'relatively small,' than for a smaller total to fall on just a few, such that their additional burden is substantial. (Temkin 2012, 67–68)

The Consolidate Additional Benefits View: in general, if additional benefits are dispersed among different people, it is better for a given total benefit to be consolidated among a few people, such that each person's additional benefit is substantial, than for a larger total benefit to be dispersed among a vastly larger number of people, so that the additional benefit any single person receives within her life is 'relatively small.' (Temkin 2012, 68)

I don't have time to repeat my arguments for these principles here, but in my book, I noted that most people, including many consequentialists, accept such principles. Such principles reflect an *anti-additive-aggregationist* approach, which opposes the simple additive-aggregationist approach of classical utilitarianism. For most people, we don't simply care about *how much* utility obtains in a given outcome, we *also* care about how that utility is *distributed* throughout the outcome, and that means that most people are *not* strictly neutral, as classical utilitarianism requires, as to *where* utility is located, with respect to people, places, and times.

Moreover, importantly, our concern about how utility is distributed does *not* merely reflect our concern for *other* distributive principles, such as equality or justice. It reflects our judgment about what is relevant to making one outcome *better* than another, *even regarding utility*. Thus, there is room for judging that though one outcome has *more* utility than another, it is still *worse* regarding utility. And this is because, for many, *one* fundamental concern is about the way in which different amounts and distributions of utility *affect people*, for better or worse, and this, I believe, reflects a *person-affecting* view, and not simply an impersonal neutralist position.

The anti-additive-aggregationist principles have wide appeal, both within, and between, lives. For example, they help explain why many find Parfit's Repugnant Conclusion repugnant (Parfit 1984, 381–390; Temkin 2012, 34–35, 37, 41–42, 324–328). Even if there is more *total utility* in Z than in A, Z's utility is dispersed across *many* lives, so that each person's life is *barely worth living*, whereas A's utility is consolidated among far fewer lives, so that each person's life is *well worth living*. Here, our ranking of A as better than Z reflects the wide person-affecting view that focuses on how the people in the two outcomes fare, and rests on the anti-additive-aggregationist position that lots of tiny benefits spread across innumerable masses don't *add up, normatively*, in the way that they would need to to make Z better than A.¹⁶

Similar results apply within lives, where most people judge a long life that includes two years of *excruciating* torture and fifteen mosquito bites per month, as *worse* than a long life that contains *no* torture but *sixteen* mosquito bites per month, even as they acknowledge that if the life were long *enough*, the *total* amount of disutility would be greater in the latter life than the former. Here, too, our anti-additive-aggregationist principles tell us that the discomfort of one extra mosquito bite per month doesn't *add up, normatively*, in the way that it would need to to outweigh two years of excruciating torture. Here, too, we are not merely concerned with the *total* amount of disutility in a life, and utterly neutral as to where, when, and to whom (in this context, which person stage) it obtains; rather, we are concerned about how the disutility is *distributed* throughout the life, and how the *person* is *affected*, for better or worse, by that distribution. The plain fact is that some distributions of vast amounts of total disutility can be benign, while some distributions of much smaller amounts of total disutility can be disastrous.

So, bearing all this in mind, return to the example depicted in Diagram 1, but this time, imagine that the outcomes don't extend infinitely, but 'only' for a billion years. How do the two outcomes compare in that case? I believe, in accordance with the spirit of the Disperse Additional Burdens and Consolidate Additional Benefits Views, that O_1 would still be *better* than O_2 , even though there is twice as *much* badness in O_1 as in O_2 , and twice as *much* goodness in O_2 as in O_1 .

How could that be? Well, by now, the answer is familiar. The question, for me, isn't merely about how much *total* goodness and badness exists, but about how the *people* are *affected* for better or worse, by the distribution of whatever goodness and badness there is. In O_1 , there will be many people with good lives for every day but one, for 100 years, many with good lives for every day but one, for 99 years and 364 days, many with good lives for every day but one, for 99 years and 363 days, many with good lives for every day but one, for 99 years and 362 days, and so on. To be sure, on the *last* day of O_1 's existence, there will be a *vast* number of people who live only one day, whose lives, for that day, will be very bad. But still, the badness for them only lasts a *single day*, and I don't believe that that badness *adds up, normatively*, across the lives of the many who will live *only one day*, in the way that it would need to outweigh the *really good lives* that would have been lived for *many* years by *many* others.

In O_2 , on the other hand, there will be many people with bad lives for every day but one, for 100 years, many with bad lives for every day but one, for 99 years and 364 days, many with bad lives for every day but one, for 99 years and 363 days, many with bad lives for every day but one, for 99 years and 362 days, and so on. To be sure, on the *last* day of O_2 's existence, there will be a *vast* number who live only one day, whose lives, for that day, will be very good. But still, the goodness for them only lasts a *single day*, and here, as before, I don't believe that that goodness *adds up, normatively*, across the lives of the many who live *only one day*, in the way that it would need to outweigh the *really bad lives* that would have been lived for *many* years by *many* others.

In sum, looking at how all the different *people* are *affected* for better or worse in each outcome, and taking account of the anti-additive-aggregationist principles that I find compelling in cases like this, as well in many other cases, including the Repugnant Conclusion, I would judge that, *even regarding utility*, Outcome One is *better* than Outcome Two. Of course, in making this judgment, I am not denying the obvious truth that Outcome Two has *more* utility than Outcome One. Of course it does! Rather, I am reflecting the appeal of person-affecting views, that pay attention not merely to *how much* utility there is in any outcome, but to how the *people* are *affected*, for better or worse, by the distribution of however much utility there is.

If my judgment about this case is correct, and I realize that not everyone will agree with me about this, then we have reason to reject the Spatial and Temporal Dominance Principles *even in finite cases*. After all, as before, Outcome Two is better than Outcome One at *every* location in space and time. At first blush, this is a somewhat surprising result. However, on reflection, I believe it is the right one.

6. The Limbo Man, the Capped Model, and other unresolved issues

In discussing the badness of death, Kamm (1993, 19, 49–55) introduces several variations of a case she calls the *Limbo Man*. Kamm suggests that the *finality* of death, its *permanence*, the *extinction* of our lives may play a special role in our explaining some of the attitudes toward death that are almost universally shared. Kamm suggests that the recognition that once our lives are over, they are over *forever*, opens the possibility that there might be reason to be concerned about when our lives begin, as well as when our lives end.

Kamm suggests, for example, that if the universe spanned a given period of time, and we were going to live our one and only life for 70 years sometime within that span, then even if one assumed that the conscious experiences of our life would be *exactly* the same on either scenario, and that all the (non-location-in-time-related) *goods* of life that we would possess would be *exactly* the same on either scenario, there might still be reason to want our 70-year life to obtain towards the *end* of the universe rather than towards the beginning.

Kamm suggests that there might be at least two related reasons for this. First, as long as we have not yet existed, the *potential* for our existing will still be there, and one might believe that there is value in such potential. Second, one might believe that it is better for us, or our lives – though *not* due to any impact on our states of consciousness or the goods that we possess during our lives – if the time during which we will never again exist is as short as possible, so that our extinction comes as late as possible. On this view, the badness of *never existing again* is distinct from, and has special significance, relative to the mere badness of *not* existing, which, of course, will also be true of us at each moment *before* we come to exist.

For Kamm, then, if someone had a choice of living a normal lifespan, filled with a given set of experiences and goods of life, or a chance of starting that life for a period of time, going into limbo for an extended period of time, and then finishing off the remainder of one's life many centuries or more later, there could be reason to do the latter, even if it were *no better* in terms of one's set of experiences and (non-time-related) goods of life. As indicated, for Kamm, this is because there might be something good about both preserving, as long as possible, the possibility of a period of one's future existence, and minimizing, as much as possible, the period during which it is true that you will *never exist again*.

In many respects, Kamm's discussion is orthogonal to my own. She is concerned about the badness of death, and the asymmetry between our attitudes towards *death*, and the period during which we won't exist that comes *after* it, and our attitudes towards *birth*, and the period during which we won't exist that comes *before* it. I am focused on civilizations, containing large populations, and my concern is not with the prospect of any given civilization coming to an *end*, but with the possibility of there being large periods of time devoid of

any civilizations at all whose members possess high-quality lives. But though Kamm's concerns are different than mine, they have a bearing on mine in several respects.

First, I have pointed out an asymmetry in my thinking about space and time. I note that the ubiquitous fear of death that has been almost universally held throughout the history of humanity, and which has inspired so much art and literature, is a *time*-related attitude. It is *not* a *space*-related attitude. Kamm's discussion takes as its starting point the commonplace that the vast majority of humans fear their own death. For many, the prospect that there will be an eternity of time after they die during which they *will never exist again* is *terrifying*. Many can't bear to even contemplate that prospect, and many others simply refuse to accept it. Indeed, throughout human history, religions have arisen to help people confront their earthly deaths, holding out the promise of eternal life.

Nothing akin to this holds regarding the many different *spaces* beyond ours that we will *never, ever, occupy!* Looking out into space, we might feel tiny and insignificant. We might even find ourselves deeply *disappointed* that we never get to explore that great unknown. But we don't look into space with utter *terror* at the realization that *we will never be there*. We have no trouble contemplating, or accepting the fact, that there may be an infinity of spaces, beyond all those that we will ever occupy. And no religions have arisen to help us cope with that reality!

Kamm's Limbo Man lives for a period, puts himself in limbo, then lives out the duration of his life at a much later period of time. In doing this, he gains nothing in terms of his experiences or the (non-time-related) goods of life, but he succeeds in significantly delaying the day when it will be true that he will no longer ever exist again. Kamm thinks that there could be reason to be this kind of Limbo Man. But Kamm doesn't consider a Limbo Man who lives in different locations, puts himself in limbo, then lives out the duration of his life at a place much further away in space than the place he was when he went into limbo. This is, I believe, no accident.

If Kamm is right, there can be rational significance to *when* we live in time, connected to the desirability of both the potential that we shall one day exist in the future, and the delay of the time after which we shall never exist again. But there seems to be no analogous rational significance to *where* in space we live. Assuming that our experiences and (non-space-related) goods would be the same, the potential of our coming to exist at one point in space seems no more valuable than the potential of our coming to exist at any other point in space, far away from the first. Nor does there seem to be any rational reason to want to start our lives at one point in space, and finish them at another, merely so as to reduce the amount of space outside the first point that will never be occupied by us.

In sum, I suggest that Kamm's discussion of our attitudes towards death – attitudes which are decidedly time related but not space related – implicitly support

my contention that there is reason to treat space and time differently for the purposes of practical reasoning. However, Kamm's discussion is of interest to me not merely for that reason, but because it raises a host of interesting questions about individual lives that might be similarly raised about large groups of lives.

In this article, I have contended that it is more important that large gaps in time be filled with many high-quality lives, than that large gaps in space be filled with high-quality lives. But even if that is so, there are a host of other, Kamm-like questions that might be raised. For example, suppose that there will only be one super-large civilization in the whole of the universe's existence, and that the very same people will exist in that civilization and all have lives of the very same high quality of existence, no matter when in time that they existed. If there would be, say, a *quintillion* (1,000,000,000,000,000,000) people, each of whom lived for 100 years with a high quality of life, would it matter if they lived later, rather than earlier, in the life of the universe? For example, if the life of the universe, and time itself, had a beginning and an end, and all of the people were going to be alive during the same 10,000-year period, would it be better if they all lived in the middle rather than the beginning, but better still if they all lived at the end?

Similarly, even if I am right that, in general, it would be good if the different 'empty' periods of time were 'filled' with high-quality lives, that still leaves many possibilities open. For example, suppose that we thought that 100 billion people were enough to 'fill' any given 10,000-year period of time. Then, on my view, it would be better if the quintillion people were spread out in time, so that there were 1000 *distinct* 10,000-year periods, each filled with 100 billion people. Still, that leaves it open whether we think it matters *how* those different periods were dispersed throughout the life of the universe.

Retaining our previous assumption that the universe, and time, had a beginning and an end, one might think that it doesn't matter. Or one might think, perhaps influenced by Kamm, that it would be best if the different periods were all bunched together, so that there was a steady run of high-quality life for 10 million years coming at the *very end* of time. Perhaps one might think that it was important that high-quality life span the entire life of the universe, but that all this required is that the very *first* 10,000-year period be populated with lots of high-quality life, and that the very *last* 10,000-year period be populated with lots of high-quality life, but that beyond that it wouldn't matter how the different 10,000-year periods of high-quality life were distributed through time, as long as they remained non-overlapping. Alternatively, perhaps one might think it best to distribute the different periods of high-quality life throughout time, so as to minimize the length of any period of time during which there would be no high-quality life!¹⁷

These questions are not intended to be exhaustive, but merely indicative of the wide range of issues that need to be considered once one acknowledges that, *ceteris paribus*, it is good if 'empty' periods of time be 'filled' with

high-quality lives. And, of course, similarly issues will arise if one believes that there is some reason to 'fill' empty locations in space with high-quality lives, even if one grants that such reasons are less weighty than the analogous ones pertinent to time. Moreover, we may or may not have the same view of these issues, depending on whether or not we think that time and space are finite or infinite.

I mention these issues, only to leave them aside. They are a reminder – as if any were needed! – of the very preliminary nature of this work.

In *Rethinking the Good*, I introduced a position I called the *Capped Model of Moral Ideals* (Temkin 2012, 328–350). On a Capped Model, in certain contexts of comparison, there may be an upper limit on how good a given outcome may be regarding any particular ideal and, similarly, an upper limit on how good an outcome can be, all things considered. I acknowledged that the Capped Model faced a host of serious objections, and that there were a large number of unresolved questions regarding it. Even so, I argued that we needed something like the Capped Model in order to capture the anti-additive-aggregationist views that underlie certain of the judgments that most people hold about how different outcomes compare. Thus, for example, in the Repugnant Conclusion, most people believe, in accordance with a Capped Model, that no matter *how many people* there may be in an outcome where everyone who exists have lives that are barely worth living, the *value* of the well-being in that outcome will never rise to the point that it outweighs the *value* of the well-being in an alternative outcome where 10 billion people exist, all of whom have extremely high-quality lives. And likewise, most people believe that the all things considered value of the former outcome will never exceed that of the latter.

I cannot repeat, here, my lengthy discussion of the Capped Model in *Rethinking the Good*. But I note that much of this article is implicitly reflecting different views as to how best to develop the Capped Model. For example, I am suggesting that we should not have a single cap for how good an outcome may be regarding well-being over the course of time, but rather that we must have separate caps for each 'substantial' period of time. Regarding space, however, I am tempted to the view that a single cap that covers all of space at any given time period might be appropriate. On the other hand, if we do have different caps for different regions of space, perhaps the levels of those caps will be lower than the levels of the caps for time, or perhaps the levels of the caps for space may vary depending on how full other spatial regions already are.

Many of the unresolved issues in this paper are connected with unresolved issues regarding the Capped Model. Correspondingly, thinking further about what to say about the kinds of cases presented in this article may help illuminate how we should think about the Capped Model. By the same token, in the Rawlsian spirit of reflective equilibrium (Rawls 1971),¹⁸ thinking further about how best to understand and interpret the Capped Model may help inform our judgments – sometimes confirming and sometimes leading us to revise our intuitions – about the sorts of cases this article considers.

Unfortunately, I cannot explore the Capped Model here. But let me point out just one way in which thinking more about the Capped Model might help illuminate our thinking about the kinds of issues this article raises. Regarding well-being, I believe that for each time period, we will need a different cap – a different upper limit – on how good an outcome can be regarding well-being, that will largely be a function of the (well-being) level of the best-off people in that outcome, and the number of people at that level. So, for example, for any given time period, the upper level of how good an outcome can be, regarding well-being, when everyone has a life that is barely worth living will be much lower than the upper level of how good an outcome can be, regarding well-being, when everyone has a life that is well worth living. More generally, I believe that the higher quality lives a population possesses, the higher cap there may be on how good an outcome can be regarding well-being.

Let us suppose that if people have a very high quality of life, k , the upper level on how good an outcome can be when everyone in a given time period has lives of quality k , is n . Let us further suppose that if people have a very high quality of life, $k + x$ where x is a positive number, the upper level on how good an outcome can be when everyone in a given time period has lives of quality $k + x$, is $n + y$. If x is a relatively small, we can assume that the difference between n and $n + y$ won't be too significant. Next, suppose that a population of 10 billion people is sufficiently great, that if 10 billion people are all at level k , the value of that outcome will be *very close* to the upper limit for that outcome, namely n . And likewise, suppose that a population of 10 billion people is sufficiently great, that if 10 billion people are all at level $k + x$, the value of that outcome will be *very close* to the upper limit for that outcome, namely $n + y$.

By hypothesis, it will then be the case that the value of the outcome in which 10 billion people are at level $k + x$, will be greater than the value of the outcome in which 10 billion people are at level k , but not *much* greater. By the same token, given the nature of the Capped Model, since, by hypothesis, a population of 10 billion is enough to get a population *very near* the upper limit for how good that outcome can be regarding well-being (depending on the levels of those 10 billion people), it follows that if 20 billion people were at level $k + x$, during the very same time period, rather than only 10 billion people, that would improve the overall value of the outcome regarding well-being *slightly*, but *only slightly*. And the same is true, of course, if 20 billion people were at level k , during the very same time period, rather than only 10 billion people. That, too, would improve the overall value of the outcome regarding well-being *slightly*, but *only slightly*. Since, by hypothesis, we have assumed that the overall value of $n + y$ is not much higher than the overall value of n , it follows that having 20 billion people at level $k + x$, during any given time period, will be not much higher than having 20 billion people at level k , during any given time period.

Now suppose that the only people living in the universe would be the 20 billion people at level k . But that they could either all live during the same time

period, in which case the total value of well-being in that outcome would be very close to n , or there could be 10 billion people living in each of two distinct time periods, each living at level k . In *that* case, there would be two distinct time periods during which the value of the well-being would be nearly n . Even if one is not a simple additive aggregationist for determining the value of an outcome – as one won't be if one adopts a Capped Model – it seems, given the considerations of this article, that there might be strong reason to attach significantly greater value to the outcome in which the 20 billion people were spread out over two distinct time periods, than to the outcome in which all 20 billion people lived during the same time period.

If this reasoning is all correct, then we might believe that there is strong reason to attach greater value to the outcome in which 20 billion people are all at level k , but spread out over two distinct time periods, than to the outcome in which all 20 billion people are at level $k + x$, but they all live during the *same* time period. But notice, we might suppose that it would be the *very same people* in the two different outcomes. In that case, we would be valuing one outcome more highly than another, because of how the well-being in that outcome is distributed *temporally*, even though the presumably better outcome would be *worse* for everyone who ever lived! This is because, in the supposedly better outcome, everyone would be at level k , while in the supposedly worse outcome, everyone would be better off, at the slightly higher level $k + x$.

I have argued at length, elsewhere, that there are a host of *impersonal* ideals – ideals that have value beyond the extent to which they are good or bad *for people* (Temkin 1993a, 1993b, 2000, 2003a, 2003b, 2003c). Such ideals open up the distinct possibility that one outcome could be *worse* than another, even though there is *no one for whom it is worse* and, similarly, that one outcome could be *better* than another, even though there is *no one for whom it is better*. Clearly, if the way in which an outcome's well-being is distributed in time or space can be relevant to our assessment of the overall value of that outcome, beyond the extent to which it affects the quality of lives (the experiences and non-time-related and non-space-related goods) of the sentient beings in that outcome, then the relevant temporal and spatial factors are *impersonal* in nature. This is an important but, on reflection, not surprising result.

Earlier in this article, I considered some examples that led me to conclude that, for *certain cases and contexts*, I favor people over times, and times over space. At the time, the reader may have been puzzled as to why I qualified my remarks in the way that I did. The preceding explains why. While there are some cases where I clearly favor how well-being is distributed across people, over how well-being is distributed across time, in some cases, such as the one I have just discussed, my judgment goes the other way. Here, as elsewhere, morality is enormously complex, and there are few, if any, simple principles that hold uniformly, and without exception or qualification, across all cases.

7. Conclusion

There is a rich tradition among game theorists, decision theorists, economists, and many philosophers, to identify rationality with the sort of impartial, neutralist, perspective commonly associated with classical utilitarianism, but with the acknowledgment that we need a wider conception of the good than classical utilitarianism endorses, and that we might need to allow for special agent-relative duties and permissions in assessing the rationality of people's beliefs, cares, and actions. On such a view, there is an important respect in which we should treat people, places, and times, the same, and be neutral with respect to all three.

I agree that there are important respects in which rationality *does* require us to treat people, places, and times the same, and to be neutral with respect to all three. But it is far from self-evident to what the substantive content of this position amounts. In this article, I have argued that there are certain cases and contexts where we ought, rationally, to treat space and time differently. I have further contended that there are certain cases and contexts where we ought, rationally, to treat people differently than space and time. I have argued that in some cases and contexts, we should give priority to people over both times and places, and to times over places. However, I have also suggested that in some cases and contexts, we should give priority to time over people.

I believe that the considerations presented in this article are plausible, so far as they go. Even so, as I have noted throughout, this article is very much a preliminary work. Accordingly, much more work needs to be done to properly assess what rationality requires of us in our treatment of people, places, and times, whether with respect to our beliefs, our cares, or our actions.

Notes

1. Throughout this article, I sometimes speak in terms of *place*, and sometimes speak in terms of *space*, depending on which sounds better linguistically in the particular context of usage. But I am using these terms interchangeably, understanding any particular place to correspond to a spatial location.
2. For the sake of simplicity, I follow Parfit (1984) in my presentation, by putting my discussion in terms of being neutral with respect to different 'people.' However, in this context, the scope of the word 'people' needn't refer only, or to all, *human* beings; it might refer to all *rational* beings or, as it does for many, to all *sentient* beings.
3. To be clear, and fair to Parfit, nothing in Parfit's argument against the self-interest theory commits him to denying this claim. To the contrary, Parfit could, and would, accept that there are *numerous* respects in which we should treat persons, places, and times differently, and he could, in fact, accept many of the claims and arguments that I shall be making in the following sections. The point is just that while Parfit has illuminated an important truth, not to be lost sight of, that in *certain respects* we should treat persons, places, and times the same, I hope to illuminate a different, but compatible, important truth, *also* not to be lost sight of, that in *certain other* respects, we should, and must, treat persons, places, and

times differently. In addition, I hope to illustrate some of the sometimes surprising respects in which this is so.

I might add that while Parfit could accept most, and perhaps all, of what follows, Sidgwick, perhaps, could not. This depends on whether or not one interprets Sidgwick to be a mental state theorist, who believes that the only sources of intrinsic value and disvalue are positive and negative conscious states, respectively. That view, combined with Sidgwick's additive-aggregationist maximizing approach, which assesses the value of an outcome solely as a simple additive function of the individual instances of intrinsic value and disvalue which obtain in that outcome, implies that one must be strictly neutral between people, places, and times, for the purposes of moral reasoning. Such a view is incompatible with most of the claims that I make in the following sections.

However, it is worth noting that it is not the maximizing structure of consequentialism, per se, that is incompatible with much of what follows, but rather the maximizing structure of consequentialism when it is combined with a particular conception of intrinsic value like that of the mental state theory (Here, I focus on the part of Sidgwick's view according to which there was always *sufficient reason*, and hence it was always practically rational, to act morally. In fact, as Parfit (2011, 6–7, 130–149) has pointed out, Sidgwick (1907) accepted the dualism of practical reasoning, according to which there was *also* always *sufficient reason*, and hence it was also always practically rational, to act self-interestedly).

I am grateful to Shelly Kagan (personal communication, October 2, 2015) for suggesting that I consider the theoretical underpinnings that might lead some people to think that we should, in fact, be strictly neutral between persons, places, and times for the purposes of practical reasoning. What the preceding discussion reveals is that the positions I shall argue for in the following sections are incompatible with mental state versions of maximizing consequentialism, or other theories of that ilk.

4. I realize that some of these assumptions will be controversial on certain interpretations of modern physics. For example, some believe that one cannot meaningfully distinguish between space and time, as the universe is composed of (inseparable) space/time points. Likewise, some believe that space and time had a beginning, perhaps at the moment of the Big Bang, and likewise, that space and time may have an end, depending on how much matter there is in the universe, and whether the Universe will eventually collapse on itself and everything, including space and time, will come to an end at a single point of singularity. In addition, some believe that time's passage is an illusion, and others, appealing to Einstein's theory of special relativity, will insist that the direction of time is relative to one's point of reference.

Still, there are a number of distinguished philosophers of science, metaphysicians, and physicists who would accept that the assumptions I am making are compatible with our best scientific views of the universe, and it is hard to deny that the assumptions I am making *might* have been true of our universe or some other universe (I am grateful to my colleague, the philosopher of physics Barry Loewer, who confirmed in an email on 16 October 2015, that the philosopher of physics, Tim Maudlin, 'thinks that time has an intrinsic direction, and he is willing to say that "time passes,"' and that some 'metaphysicians think that time "moves" in a more robust and non-metaphorical sense ... [including] Dean Zimmerman ... CD Broad, Michael Tooley, [and] Tim Williamson.' Loewer also noted that 'two

physicists that are in the “time really moves” camp are Lee Smolin and George Ellis.).

So, one way of interpreting the following arguments is that they may give us pause for treating space and time the same even on our best current scientific understanding of the universe, and would give us reason for treating space and time differently in any universe where something like a Newtonian conception of space and time was true. Accordingly, we should be wary of any a priori arguments in support of the claim that rationality requires that we treat space and time the same for the purposes of practical reasoning.

5. Scheffler (2013) has argued that having descendants who will help realize some of our deepest hopes, projects, or ideals helps to give our lives value and meaning that they otherwise would lack. Scheffler’s views are entirely compatible with my own, and I am happy to accept them. But they point to *other* reasons why one might be more concerned about the future than about what happens elsewhere in space than those I am trying to illuminate here. As my example makes plain, I believe that even if the future civilizations were unrelated to our own, and would do *nothing* to further *our* particular hopes, projects, and ideals, I *still* believe that there would be strong reason to ensure that such civilizations would exist if they would have high-quality lives. In addition, I believe that such reasons would be stronger than any we would have to ensure, were it possible, that such civilizations obtain elsewhere in space contemporaneous with our own.

Similarly, Jeff McMahan (personal communication, October 2, 2015) suggested a variety of considerations that might lead us, in general, to give greater weight to there being high-quality sentient lives existing in the future, than to there being high-quality sentient lives existing elsewhere in space. According to McMahan, these might include views we have about the importance of the preservation of value, views about the importance of progress, and views about the importance of greater diversity of experiences. My response to McMahan is threefold.

First, as with what I said about Scheffler’s view, I don’t regard my position as incompatible with McMahan’s. Depending on the details of the case, there could be more than one reason for valuing the existence of *future* civilizations over the existence of contemporaneous civilizations elsewhere in space. But second, in my examples, I wasn’t, in fact, assuming that there was greater diversity of experiences over time than across space, nor was I assuming that there would be *progress* between our current civilization and the future, unrelated, civilizations. Thus, my views about such cases weren’t, in fact, turning on such factors. Moreover, importantly, I note that the notions of *preservation* of value, and *progress*, have a temporal dimension built in to them, but not a spatial dimension. So, McMahan’s suggestions regarding those factors would, if correct, not be a *rival* to my own, but rather a further elucidation of some of the *reasons* why we should treat space and time differently for the purposes of practical reasoning.

Finally, Jonathan Weisberg (Q&A, *Belief, Action, and Rationality Over Time* Workshop, University of Wisconsin-Madison, September 6, 2015) suggested that our intuitions about such cases might reflect the psychological phenomenon of our ‘not wanting the story to end’; perhaps traceable to our early childhood when our parents read us bedtime stories and we didn’t want the story to end, since when it did, our parents would leave, and we would be left alone in the scary dark! It is always difficult to *prove* that such a deflationary account plays *no* role in our intuitions about such cases, but when I think hard about such cases,

I don't believe that the best, or main, explanation of my intuitions about them lies in the sort of deflationary account that Weisberg suggests.

This is further supported by cases that I will consider later, where I have similar views about the greater importance of filling *past* temporal regions, or *gaps* between present and future temporal regions, relative to the importance of filling other contemporaneous spatial regions, or 'empty' gaps between different regions of space that are occupied by high-quality sentient lives. In such cases, it appears that our propensity to 'not want the story to end' would have no explanatory role to play as to why we have the intuitions we do.

6. Thoma (2015) notes that there are two different possible ways of 'filling' space with sentient life, and wonders which one I have in mind. On one, we add new sentient beings to locations in space that were previously empty. On the other, we make previously existing sentient beings *larger*, so that they occupy more space, including some previously unoccupied space.

I had the former notion in mind, but Thoma's observation points to another issue that may have a bearing on how we should regard space and time for the purposes of practical reasoning. It is clear that my thoughts occupy time – each thought begins at one point in time, ends at another, and spans a given time period. It is much less clear that my thoughts exist in space.

As has long been observed, unlike physical objects that clearly do exist in space, thoughts lack extension, shape, or mass. To be sure, I think of my thought as occurring 'within my head,' and my head is a material substance that occupies a given region of space; but does my thought occupy the region of space from ear to ear and from chin to scalp? That sounds odd, and not quite right. Does the thought 'I am hungry' occupy *more* space, as one grows from toddler, to teenager, to adult? That doesn't seem right, either. Similarly, if I were given extraordinary growth hormones, so that I ballooned up to the size of a planet, I don't think that my thought 'I am hungry' would then be planet-sized, occupying the vast spatial region that my body would then occupy.

If this is right, then we see another way that we have to treat space and time differently, for the purposes of practical reasoning. Regarding contiguous future *temporal* regions, we could result in a net gain in the total amount of time that was filled with high-quality sentient life in either of two ways; either we could 'fill' it with new sentient beings who would have high-quality lives, or we could, in principle, 'fill' it by extending our own lifespans. But regarding contiguous *spatial* regions, it seems that we could result in a net gain in the total amount of space that was filled with high-quality sentient life in only one way, by 'filling' it with new sentient beings who would have high-quality lives. Were we to 'fill' such contiguous spaces either by moving into them, or by making ourselves larger, it seems that either way there would be no net gain in how much high quality sentient life obtained in the spatial regions in question.

Finally, this discussion makes plain that when I talk of 'high quality sentient life,' I am referring to the well-being of sentient beings that is related to *consciousness*. It also suggests that talk of 'filling' a region of space with high-quality sentient life is metaphorical, in a way that talk of 'filling' a temporal region with high-quality sentient life need not be. I am grateful to Thoma's observation and query for leading me to consider and clarify my view of these matters.

7. Thoma (2015, 10) questions whether I can legitimately conclude from my examples that 'filling space is less important than filling time.' Thoma contends that such a conclusion is indefensible in the absence of a conversion scale that would tell us what length of time is equivalent to what area in space. But Thoma

is skeptical whether such a conversion scale is even intelligible. What would it mean, Thoma wonders, to contend that one acre was equivalent to a given amount of time?

Thoma raises an interesting question. But I'm not sure how much I should be troubled by it. The question *isn't* whether, for example, a galaxy in space is equal to 100 billion years of time. Rather, the question concerns the extent, if any, to which adding billions and billions of high-quality lives to an otherwise 'empty' galaxy (one devoid of sentient life) would significantly improve the outcome, if there were *already* billions and billions of high-quality lives being lived *elsewhere* in the universe *at the same time*; vs. the extent, if any, to which adding billions and billions of high quality lives to an otherwise 'empty' 100 billion years would significantly improve the outcome if it is true that there *have been* or *will be* billions and billions of high-quality lives lived *at different times* (perhaps, to keep the situations analogous, at the very same locations to which the 'extra' lives would be added).

I don't believe that we need an answer to the former question – Thoma's question – in order to answer the latter questions – which are mine. And I believe that the answers that we get to the latter questions, and others like them, support my claim that 'it is more important for high quality life to be dispersed across time rather than across space' – at least if that claim is charitably, and properly, interpreted.

Consider someone who claims that freedom is non-instrumentally valuable; that is, valuable *itself*, beyond the extent to which it promotes other valuable ideals. Suppose that person also believes that pleasure is non-instrumentally valuable, and so believes that eating ice cream is valuable, at least for her, in virtue of the pleasure that it gives her. Echoing Thoma, one might wonder whether there is a conversion scale between the value of freedom and the value of pleasure. What would it mean, one might wonder, to contend that eating a gallon of ice cream was equivalent to a given amount of freedom? Yet, admitting that the question is interesting, and even that it raises a host of deep and important issues, I still think it could be perfectly intelligible, and even true, to assert that freedom was more valuable than eating ice cream!

Note, if we don't believe that the value of freedom is *lexically prior* to the value of pleasure – and I don't, as, setting health issues aside, I would *gladly* trade a tiny infringement on my freedom, for a lifetime of ice cream indulgence! – we may need to say a bit more to convey exactly what we have in mind in saying that freedom is more valuable than eating ice cream. But I trust that this could be done. For example, we might just mean that a lifetime devoid of ice cream eating would be better than a lifetime devoid of freedom. But the point is that we might be able to perceive the truth of such a claim, even in the absence of having a conversion scale of the sort to which Thoma is referring.

Similarly, I might not have a conversion scale between areas and times, and yet I might be correct in claiming that 'it is more important for high quality life to be dispersed across time rather than across space.' The content of my claim can be derived from the examples I presented in support of it, though admittedly, perhaps only imprecisely. That content is, I believe, intelligible and defensible, and does not require that filling 'empty' times with high-quality life has *lexical priority* over filling 'empty' spaces.

There is, of course, much more that could be said in response to Thoma's worry, but, given my space constraints, I hope to have said enough for this article's purposes.

8. Indeed, by some standards of measuring how much of each atom is 'occupied' by material entities, an atom is 99.999999999999% empty space. See, <http://physics.stackexchange.com/questions/39143/percentage-of-water-that-is-void-or-empty-space>.
9. Shelly Kagan (personal communication, October 2, 2015) wonders how far I would push the idea of the importance of filling empty times with high-quality sentient life. Suppose, for example, that there was a planet teeming with high-quality life, such that for 12 hours each day everyone was awake, and experiencing high levels of consciousness, but that for another 12 hours each day, everyone fell into a deep, dreamless, sleep, utterly devoid of conscious experiences. Would I think it a 'great tragedy' that there was so much 'empty' time over the course of each year? Would the world be better if half of the people had their sleep cycles shifted, so that at every moment, there were always large numbers of people experiencing high levels of consciousness? What if there were only five minutes each day, or five seconds each day, where everyone was completely unconscious? Would the outcome really be better, even in *such* cases, if people's sleep cycles were shifted so as to fill *every* available period of time with high levels of consciousness? Kagan finds this hard to believe.

I confess that I am tempted to the view that, if possible, it *would* always be better, other things equal, to fill *any* empty period of time, no matter how short, with many instances of high levels of consciousness; though, obviously, the shorter the period is, the less important it would be. But I'm not committed to this view. What I *am* committed to is the thought that it is bad if there are significantly lengthy periods of time – leaving it open for now whether we understand that to be on the scale of a decade, century, millennium, or eon – which are utterly devoid of high-quality life. It is better, I think, if there are large numbers of high levels of consciousness obtaining within, and spread throughout, each such period. But whether it is important that there be *no gaps at all* of high-level consciousness *within* each period, or exactly how small such gaps would have to be for them to be unobjectionable, are questions about which I have no firm views or commitments.

10. This example is a variation of an infinite one suggested to me by Jeff McMahan (after my lecture at the Oxford Moral Theory Seminar, May 18, 2015) in defense of the view in question. McMahan noted that he thought that a universe where, for all of eternity, there were billions of people in extreme agony, and no one else existed, would be worse than a universe involving an infinite number of planets, each populated by billions of people in extreme agony, and no one else existed, if, in the latter case, all of the universe's agony obtained within a single short period, such as 30 or 100 years. I agree.

Frances Kamm (personal communication, October 2, 2015) indicated that she might value a universe where lots of conscious life existed for all of eternity, over a universe where an infinite amount of conscious life existed for only 30 or 100 years, even if all of the conscious life was suffering greatly. But when pressed, it became clear that Kamm only thought this on the assumption that, all things considered, there was value to the existence of conscious life, *even if* that life involved unrelenting suffering. McMahan and I were assuming differently in making our claims. But if we agreed with Kamm, then our cases would be similar to my initial cases, where the argument purported to show that it was more important to fill time than space with conscious beings if, all things considered, the lives of those conscious beings were valuable.

11. There are at least two reasons one might think this rational. First, following Nagel (1986), one might believe that there are different perspectives on the world that one might rationally take, each of which reflects its own reality. In particular, one can take a (more or less) 'objective perspective,' which reflects the world as it is, from the 'outside,' as it were. In addition, one can take a (more or less) 'subjective perspective,' which reflects the world as it is, from the 'inside,' as it were. Even if one grants that from a certain objective perspective, there would be nothing to choose between the two alternatives, it seems undeniable that considering the *lived life* as it would be *experienced from the inside*, there would be every reason to delay the timing of one's entrance into Hell as long as possible! On Nagel's view, both perspectives reflect reality, and both are sources of genuine reasons to believe, care, or act. A fortiori, on Nagel's view, it would be perfectly rational to delay one's entry into Hell.
- Second, one might grant that there are various perspectives and theories which would regard the life with the delayed entrance into Hell and the life where one enters Hell immediately as *equally bad*. But surely there are *some* perspectives and theories that would regard the life with the delayed entry into Hell as *better* than the life that begins in Hell immediately, while there would appear to be *no* perspectives or theories that would rank the life that begins in Hell immediately as *better* than the life with the delayed entry into Hell.
- Accordingly, even if we attach *most* credence in the theories that regard both lives as equally good, as long as we attach *some* credence in the theories that regard the delayed entry life as best, and *no* credence in the theories that regard the immediate entry life as best, then from the standpoint of *practical rationality*, we have everything to gain and nothing to lose by following the theories that favor the delayed life. This follows from simple dominance-type reasoning, since if, in fact, a theory that regards the two theories as equally good is true, then we won't be any *worse off* if we pick the delayed entry life, while if, in fact, a theory that regards the delayed life as best is true, we will be better off if we follow it, and worse off if we don't! A powerful argument in support of the *practical rationality* of choosing in accordance with the line of reasoning suggested is presented and defended in Ross (2006). Ross's argument is discussed and employed at numerous points in Temkin (2012, 35–36, 40–41, 125–127, 171–173, 261–262, 443–445).
12. Note, our previous discussion would support a similar, though possibly weaker, conclusion for finite cases. If it is more important to fill temporal locations that are devoid of high-quality sentient life than spatial locations that are devoid of such life, it would be more important to shift the *temporal* locations of some people in an overcrowded world to some point in the future that would otherwise be devoid of high-quality life (say, via suspended animation, if that were possible), than it would be to merely shift the *spatial* location of the people in question by sending them to another planet that would otherwise be devoid of such life (say, via teletransportation, if that were possible). For further discussion of this position, see Section 6.
13. The following case is my own, but it was sparked by an example I first heard from John Broome, many years ago, which he called 'Expanding Heaven and Expanding Hell.' Broome credited his example to Cain (1995). Although my views about this topic were arrived at independently, other philosophers have developed similar arguments in order to make similar points. See, for example, Vallentyne (1993), Cain (1995), Lauwers (1997), Vallentyne and Kagan (1997), Machina (2000), Lauwers and Vallentyne (2004), Bostrom (2011), and Campbell (2015).

Interestingly, while Cain (1995) uses an example similar to mine to arrive at the same conclusion that I do regarding the relative status of Personal and Temporal Dominance Principles for certain cases and contexts, Campbell (2015) produces a series of ingenious examples in order to show that, depending on one's theory of personal identity, there will be other cases where the relative status of Personal and Temporal Dominance Principles would be the reverse of what Cain and I argue for. I don't happen to favor the reductionist view of personal identity that would lead to Campbell's results, but many do, and for those who do, Campbell's arguments are quite compelling.

14. As recognized in note 13, Campbell (2015) has shown that, depending on one's view of personal identity, there may be cases where the judgment yielded by the Personal Dominance Principle seems false. Moreover, in Section 6, I will present other cases where the Personal Dominance Principle seems false, that don't depend on one's views about personal identity. Thus, on reflection, I believe that the Personal Dominance Principle needs to be revised, or limited in scope, and the same is true of the other dominance principles.
15. For reasons of the sort adduced in note 11, assuming that we give *some* credence to thinking that the Personal Dominance Principle applies in such a case, then for the purposes of practical reasoning, we should choose *as if* that is the correct theory, even if in fact we give more credence, and even much more, to the Impersonal Neutralist View. This is because, for Pareto-like reasons, in *this* kind of case, we have everything to gain and nothing to lose by following the recommendation of the Personal Dominance Principle. See Ross (2006), and Temkin (2012, 35–36, 40–41, 125–127, 171–173, 261–262, 443–445).
16. My use of the terms 'narrow person-affecting view' and 'wide person-affecting view' vary in certain important respects from Parfit's use of those terms, but in ways that need not concern us here. See Parfit (1984, 393–401; Temkin 416–45).
17. I am grateful to Frances Kamm (personal communication, October 2, 2015), for leading me to see that there were a slew of issues of this sort that ultimately need to be considered and resolved.
18. The methodological approach of seeking 'reflective equilibrium,' famously championed by Rawls, was, as Rawls himself acknowledges, previously employed by Sidgwick (1907).

Notes on contributor

Larry S Temkin is a distinguished professor and chair of the Department of Philosophy at Rutgers University. He is a specialist in ethics, and social and political philosophy, and the author of *Inequality* (Oxford University Press, 1993) and *Rethinking the Good: Moral Ideals, and the Nature of Practical Reasoning* (Oxford University Press, 2012).

ORCID

Larry S. Temkin  <http://orcid.org/0000-0002-1194-4379>

References

- Bostrom, Nick. 2011. "Infinite Ethics." *Analysis and Metaphysics* 10: 9–59.
- Cain, James. 1995. "Infinite Utility." *Australasian Journal of Philosophy* 73: 401–404.
- Campbell, Tim. 2015. "Personal Ontology and Bioethics." PhD diss., Rutgers University.
- Kamm, Frances. 1993. *Morality, Mortality Vol I: Death and Whom to Save from It*. New York: Oxford University Press.
- Lauwers, Luc. 1997. "Infinite Utility: Insisting on Strong Monotonicity." *Australasian Journal of Philosophy* 75: 222–233.
- Lauwers, Luc, and Peter Vallentyne. 2004. "Infinite Utilitarianism: More is Always Better." *Economics and Philosophy* 20 (2): 307–330.
- Machina, Mark. 2000. "Barrett and Arntzenius's Infinite Decision Puzzle." *Theory and Decision* 49 (3): 291–295.
- Nagel, Thomas. 1986. *The View from Nowhere*. New York, NY: Oxford University Press.
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Oxford University Press.
- Parfit, Derek. 2011. *On What Matters*. 1 vol. Oxford: Oxford University Press.
- Rawls, John. 1971. *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Ross, Jacob. 2006. "Rejecting Ethical Deflationism." *Ethics* 116: 742–768.
- Scheffler, Samuel. 2013. *Death and the Afterlife*. New York: Oxford University Press.
- Sidgwick, Henry. 1907. *The Method of Ethics*. 7th ed. London: Macmillan.
- Temkin, Larry. 1993a. *Inequality*. New York: Oxford University Press.
- Temkin, Larry. 1993b. "Harmful Goods, Harmless Bads." In *Value, Welfare, and Morality*, edited by R. G. Frey and Christopher Morris, 290–324. Cambridge: Cambridge University Press.
- Temkin, Larry. 2000. "Equality, Priority, and the Levelling down Objection." In *The Ideal of Equality*, edited by Matthew Clayton and Andrew Williams, 126–161. London: Macmillan.
- Temkin, Larry. 2003a. "Equality, Priority or What?" *Economics and Philosophy* 19: 61–87.
- Temkin, Larry. 2003b. "Egalitarianism Defended." *Ethics* 113: 764–782.
- Temkin, Larry. 2003c. "Personal Versus Impersonal Principles: Reconsidering the Slogan." *Theoria* 69: 20–30.
- Temkin, Larry. 2012. *Rethinking the Good Moral Ideals and the Nature of Practical Reasoning*. New York: Oxford University Press.
- Thoma, Johanna. 2015. "Larry Temkin: Rationality with Respect to People, Places and Times." Paper presented at the Belief, Action, and Rationality over Time Workshop, University of Wisconsin-Madison, Madison, WI, September 6.
- Vallentyne, Peter. 1993. "Utilitarianism and Infinite Utility." *Australasian Journal of Philosophy* 71: 212–217.
- Vallentyne, Peter, and Shelly Kagan. 1997. "Infinite Value and Finitely Additive Value Theory." *The Journal of Philosophy* 94 (1): 5–26.
- Williams, Bernard. 1981a. "Persons, Character and Morality." Chap. 1 in *Moral Luck*. Cambridge: Cambridge University Press.
- Williams, Bernard. 1981b. "Internal and External Reasons." Chap. 8 in *Moral Luck*. Cambridge: Cambridge University Press.
- Williams, Bernard. 1985. *Ethics and the Limits of Philosophy*. Cambridge, MA: Harvard University Press.