

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/93943>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

Past Killings and Proportionality in War*

When war starts, things often don't go as planned, so we make new plans; either plans to continue or plans to stop. Our question is: what are the implications of the fact that things have gone awry for our new plans? Obviously, the fact that things haven't gone as planned may have epistemic significance. War is more trial and error than some would like to admit. And our earlier acts can change the facts – they can make it harder for us to achieve our aims, for example by strengthening the resolve of our opponents, or by making it more likely that insurgent forces will disrupt our plans, or by demoralizing our forces so that they are less likely to succeed if the war continues.

Often, a country that commits itself to war finds it hard to extricate itself, for both practical and psychological reasons. Withdrawing from a war may make a country seem weak, and that may make it more vulnerable to future attacks; it is difficult for politicians who have supported a war to concede that they have wrongly ordered their own combatants to face the burdens of combat; and politicians and military leaders are motivated to ensure that the losses that combatants and others have suffered are not in vain. These motivations and others can lead wars to go on longer than they ought to, even when it is clear that the goals of war are

* This article builds on some arguments that I developed in the course of discussing C Fabre 'War Exit' (2015) 125 *Ethics* 631 on PEA Soup: <http://peasoup.typepad.com/peasoup/2015/05/ethics-discussions-at-pea-soup-cecile-fabres-war-exit-with-critical-precis-by-helen-frowe.html>. I am grateful to Cécile for helping me to develop my views. I am also grateful to the war discussion group at Oxford for their thoughts about an earlier draft. Thanks to Seth Lazar for sharing his unpublished work on this issue, and to Jeff McMahan and Darrell Moellendorf for discussion of the topic. I am especially grateful to the Leverhulme Trust for a Major Research Fellowship that afforded me the time to work on this article.

insufficiently important to justify continuing to fight. WWI and the US war in Vietnam are perhaps the most obvious examples.

The question I address here is concerned with the moral difference that losses inflicted in war make to the decision to continue: whether the harm that a country causes early in the war makes a difference as such to what it is now permitted to do – whether it makes a difference not simply because of the difference it makes to its evidence, or because the facts have changed. I argue that one way in which the past has been thought significant in this way is either insignificant, or not very significant: the fact that some innocent people have been killed as a side-effect of a war does not affect forward-looking wide proportionality calculations as such, or if it does it does not affect them much.

Whether harm-causing conduct is widely proportionate depends on the relationship between the good that the harm-causing conduct achieves, or is reasonably expected to achieve, and the harm that the conduct inflicts, where those who are harmed are not liable to suffer that harm. If they are not so liable they have a right not to be harmed. Harm-causing conduct can nevertheless sometimes be justified because the good that will be done by the conduct is sufficiently important to justify infringing the rights of others. When it is, that conduct is widely proportionate. Wide-proportionality contrasts with narrow proportionality, which is concerned with harming liable people.¹

To clarify the issue I am concerned with, I will focus on variations on this case:

Early Losses. Country X goes to war with country Y to save 50000 people who will otherwise be killed by Y's officials. According to X's evidence at the time of going to war, t1, doing this is expected to kill 10000 innocent people. That number (I stipulate) is small enough that going to war is proportionate, but only just. It would be disproportionate to kill 12000 innocent people to save

¹ Jeff McMahan developed the distinction between wide proportionality and narrow proportionality. See, especially, *Killing in War* (Oxford: OUP, 2009) 20-1.

the 50000. The war begins at t_1 , but unfortunately things do not go as planned. Early battles are lost and 10000 innocents are killed as a side-effect. At t_2 , X does a new calculation. According to its evidence, the 50000 can still be saved, but some more innocents will be killed.

I will treat country X as the relevant agent for evaluation, though nothing turns on this. Suppose that given other facts about the war, whether continuing to fight is permissible depends on whether doing so is widely disproportionate. If the future losses are sufficiently small for this to be so, continuing to fight is permissible; if it is too large, continuing to fight is wrong. Our question is whether the fact that X has already killed 10000 makes a difference as such to wide proportionality judgements.

Here are three scenarios:

Early Losses 1: Saving the 50000 is expected to result in 2000 more deaths.

Early Losses 2: Saving the 50000 is expected to result in 10000 more deaths.

Early Losses 3: Saving the 50000 is expected to result in 12000 more deaths.

Here are four views:

Quota: If X's evidence warrants the belief that the total number of deaths that will be caused to save the 50000, including those killed in the early battles and those still to be killed, will make the war fact-relative widely disproportionate as a whole if it continues to fight, X ought not to continue to fight at t_2 .

Discount: The fact that X has caused 10000 deaths in the effort to save the 50000 counts against causing further deaths at t_2 , and can make it widely disproportionate to continue fighting, but each of these deaths count for less against continuing to fight than each prospective death.

Prospect: The fact that X has caused 10000 deaths in the effort to save the 50000 does not count at all either for or against continuing to fight at t2.

Addition: The fact that X has caused 10000 deaths in the effort to save the 50000 counts in favour of continuing to fight at t2, and can make it proportionate to kill more people to save the 50000 than was the case at t1.

As we will see, fully specifying these views (other than *Prospect*) requires us to be careful about which past losses affect future decisions. For example, does it depend on whether these losses are caused by this country; or in this war; or in service of saving this 50000? Let us leave these complications aside for the moment.

If wide proportionality is all that is at stake, *Quota* implies that continuing to fight is wrong in all three *Early Losses* cases. There is a quota of deaths that X may cause to save the 50000. It is wrong for X to exceed this quota, as it will in all three cases. *Discount* need not have this implication. It might imply that continuing to fight is permissible in *Early Losses 1* but wrong in *Early Losses 2* and *Early Losses 3*. *Prospect* implies that continuing to fight is permissible in both *Early Losses 1* and *Early Losses 2*, but wrong in *Early Losses 3*. *Addition* implies that continuing to fight is permissible in *Early Losses 1* and *Early Losses 2*, and even perhaps in *Early Losses 3*.

Whilst I restrict my discussion to the context of war, it has wider application. It is part of a more general theory of how the costs of our past actions affect what we ought to do. That more general question includes the rationality of continuing to pursue a course of conduct that has proved costly to me – what economists refer to as the problem of sunk costs. The discussion in this paper, though, is restricted to the question of the effect of past side-effect costs of our actions on the proportionality of inflicting future side-effect costs on others. That is not quite the issue of sunk costs, as the relevant costs are not costs borne by the agent deciding what to do, but are borne by others. Even the issue I focus on here is important beyond the context of war, though. For example, it is permissible to build public buildings only if the disruption to local residents is not too great. If the disruption caused in the early stages of building is greater than expected, and less progress is made than expected,

what difference does past disruption make to the permissibility of continuing to build? But addressing the problem is especially important for the context of war, where life and death is at stake.

Section I shows both that the two main attractions of *Quota* are illusory. Section II offers a decisive argument against *Quota*. Section III suggests that even if lost lives count against continuing to fight, they cannot count against doing so in the same way as prospective losses. Section IV responds to an argument that might be offered in favour of *Discount*: one that suggests that past harms are clearly relevant in intrapersonal cases, so they must be relevant in interpersonal cases. Section V rejects further arguments that the rights of the 50000 to be saved are weakened, or the duties of X not to kill are strengthened, by past losses. Section VI rejects the two best arguments that I can find in favour of *Addition*. Overall, I support *Prospect*.

I. The Illusory Appeal of Quota

I am most confident that *Quota* is wrong. In this section I show that three reasons given for *Quota* do not support it. In the next section I offer a decisive argument against it.

i) *Evaluating the War as a Whole.*

One argument for *Quota* is that it seems that if X continues to fight at t2 in *Early Losses*, the war as a whole will be disproportionate, and thus X will have acted wrongly overall. Some think that only *Quota* can adequately explain this judgement. Darrell Moellendorf expresses this view thus:

A central question in the morality of continuing the war in, and occupation of, Afghanistan has been whether the realization of the cause is worth the total moral costs. That seems like a meaningful question...If we were to employ a conception of proportionality that looks only forward, any such criticism would be fundamentally confused...Such a conception of proportionality is

unable to make sense of an important part of the current practice of arguing about war. It, therefore, does a grave disservice to our moral understanding.²

It may be true that we need a way of capturing the idea that the war as a whole is disproportionate, but friends of *Prospect* have a way.

To see this I rely on a distinction between kinds of wrongness that Derek Parfit has made familiar: the distinction between fact-relative and evidence-relative wrongness.³ In the fact-relative sense, whether a person has acted wrongly depends solely on the facts, regardless of whether the person is aware of those facts, or has evidence of them.⁴ In the evidence-relative sense, it depends on the evidence available to that person. For example, suppose that I serve you a drink which has poison in it, but I have no evidence that there is poison in it. What I do is fact-relative wrong but evidence-relative permissible. In contrast, if I give you a drink which does not have poison in it, but I have strong evidence that it does, my act is fact-relative permissible but evidence-relative wrong.

Now suppose that *Prospect* is right, and X is permitted to continue fighting in *Early Losses 1*. It does so, kills 2000 as expected, and saves the 50000. X's total set of acts from t1 to the saving of the 50000, are wrong in the fact-relative sense, because

² See D Moellendorf 'Two Doctrines of *Jus ex Bello*' (2015) 125 *Ethics* 653, 667. For a similar idea, see C Fabre 'War Exit' (2015) 125 *Ethics* 631, 637.

³ See, especially, D Parfit *On What Matters vol.1* (Oxford: OUP, 2011) ch.7.

⁴ Parfit suggests that an act is fact relative wrong 'just when this act would be wrong in the ordinary sense if we knew all of the relevant facts' (*On What Matters vo.1*, 150). The counterfactual test that Parfit proposes is best understood as an imperfect heuristic device for fact-relative wrongness rather than an account of what it is, because knowledge of the facts may be morally important as such. For example, on the best sense of fact-relative wrongness, it might be fact-relative wrong for me to declare that I am knowledgeable when I am not, but it would not be wrong for me to declare that I am knowledgeable were I to know all of the relevant facts. Nothing turns on this refinement for our purposes though.

this total set of acts has saved 50000 at the cost of 12000, which is widely disproportionate. As these acts of war constitute the war, the war as a whole is disproportionate, and because of that the war as a whole is wrongful in the fact-relative sense. We need not claim that X acts wrongly at t2 for this judgement to be warranted.

It might be argued that X should stop fighting because if it continues the conjunction of its acts will be fact-relative wrong. But this conclusion is too quick. The conjunction of X's acts is fact-relative wrong whatever X does at t2. If X kills 10000 between t1 and t2, and stops fighting at t2, X has killed 10000 people to save no one at all. This set of acts, taken together, would also be fact-relative wrong, for at t1 it would obviously have been wrong for X to kill 10000 people in a way that would not save anyone. X has decisive evidence that if it stops now, it will have killed 10000 people and saved no one. So at t2 it has decisive evidence that its acts of killing and then stopping are together wrong.

Once we see that X's conjunction of acts at t1 and t2 is fact-relative wrong whatever it does, some of the appeal of *Quota* is lost.⁵ Some of its appeal is due to the judgement that rejecting it seems to involve giving up on our initial proportionality assessment at t1, or treating that assessment as morally unimportant. But at the same time, we must make room for an equally important assessment – that killing 10000 people to save no one is gravely morally wrong. The impression that *Quota* has an advantage because it can explain the first assessment is dispelled once we see that it fails to explain the second.

ii) *Protecting Interests and Respecting Rights*

A second idea is that only *Quota* adequately protects the interests of those who will be killed after t2. Cécile Fabre argues as follows. The interests of these people would have protected them at t1 had X known the facts: they would have counted

⁵ For a related argument, see J McMahan 'Proportionality and Time' (2015) 125 *Ethics* 696, 706-7.

decisively against saving the 50000 at t1. If X is permitted to proceed at t2, they are no longer protected. She concludes that only *Quota* provides an adequate constraint against killing.⁶

In response, it is true that the 2000 lack a protection at t2 that they would have had at t1 had X known all the facts. The question is whether this is troubling. Whether our interests are sufficiently important to make it wrong for others to set those interests back often depends on other facts, such as whether the act that sets back our interests also sets back other people's interests. For example, whether my interest in not being killed protects me from being killed as a side-effect of saving 50000 people depends on whether I am a member of a group of people that is sufficiently large to make the saving of the 50000 disproportionate. If that group is sufficiently large, my interest protects me. If it is not, it doesn't.

In *Early Losses 1* these facts have changed. At t1, the 2000 were part of a larger group of 12000 people whose interests would all have been set back by the acts that save the 50000, and because they were part of this large group their interests would have protected them from being killed had X known all of the facts. But at t2 they are no longer in this group. The 10000 are now dead, so killing the 2000 at t2 no longer sets back the interests of the 10000. So the interests of the 2000 no longer protect them.

Perhaps it might be argued that in proceeding, X would fail to respect the rights of the 2000. Their rights would have protected them from being killed were X to have known all the facts at t1. The proper way to respect their rights is to refrain from acting at t2 in a way that would have constituted a rights violation at t1.

In response, respecting the rights of the 2000 involves giving the interests of the 2000 the proper weight in our deliberations. *Prospect* is compatible with X doing this at both t1 and t2. As X's evidence justified going to war at t1, *ex hypothesi*, X has not failed to respect these rights at this stage. At t2, X weighs the interests of the 2000 and the 50000 who can still be saved. If *Prospect* is right, the lives of the 10000 make no difference to how these interests should be weighed. But that in no way

⁶ See Fabre 'War Exit' 637.

diminishes the importance of the rights of the 2000 in assessing what X should do. It is just that in X's new circumstances, X has decisive evidence that they are outweighed by the interests of the 50000.

It is true that the rights of the 2000 would have been sufficiently important to make it wrong for X to go to war at t1 *had X known the facts*. But X has not violated the rights of the 2000 between t1 and t2; these people have not yet been affected at all. If X kills them, their deaths will make the war as a whole fact-relative wrong, because too many people were killed to save the 50000, and they were part of the group of people whose deaths make this true. But that does not imply that X ought to desist, for as I have already noted, X's acts together will also be wrong in the fact-relative sense if X desists. That will be so because the rights of the 10000 will have been violated in the fact-relative sense.

It is true that X's fact-relative wrongful acts at t1 makes it the case that the rights of the 2000 do not protect them. But it is hard to see why the fact that X has acted wrongly towards the 10000 at t1 should count in favour of protecting the 2000. Compare this 2000 people with another group of 2000, who could be killed to save 50000 other people, but who were never part of a larger group that made saving the 50000 fact-relative wrong. Why should the latter group receive less protection from being killed than the former group simply because, as a matter of luck, the rights of the former group made it true that X acted wrongly in the fact-relative sense at t1?

Perhaps it will be argued that if X proceeds it will have violated more rights than if it desists. If X fights at t2 it will have violated 12000 rights. If it desists it will have violated 10000 rights. It is plausible that in a disproportionate war, the right to life of each person whose death makes the war disproportionate is violated. But the view that this consideration favours stopping gives no weight at all to the interests of the 50000, who will be killed if X does not continue. Their interests counterbalance the rights of those who will be killed.

Some may argue that as the 50000 would not have had a right to be saved had X known all the facts at t1, their rights are not engaged at t2. But the idea that this argument rests on cannot be right. The reason why the 50000 would not have had a right to be saved is because of the 10000 who have been killed between t1 and t2. The

fact that the deaths of the 10000 would have made it wrong for X to save the 50000 cannot negate the significance of the interests of the 50000 at t2 altogether.

To see this, consider:

No Cost: As *Early Losses* except 12000 are killed between t1 and t2, but the 50000 can be saved at no cost to anyone.

I stipulated that it is disproportionate to kill 12000 to save 50000. Had X known all of the facts at t1, it would thus have been wrong for X to act. But it is surely wrong not to save the 50000 from death at t2 where this can be done at no cost. And failing to do this would surely wrong the 50000: they have a right to be saved where this can be done at no cost. Thus the fact that they would not have had a right to be saved at t1 does not completely undermine their right to be saved at t2. It might be argued that the rights of the 50000 are significantly weakened, even if they are not lost altogether. But this is an argument for *Discount*, not *Quota*. We will consider it below.

iii) *Iteration*

A second familiar argument that has been offered to support *Quota* concerns iteration. To see the problem that this argument is concerned with, suppose that in *Early Losses 2*, things again don't go as planned at t2. 10000 more people are killed at the next stage of fighting, and these battles are also lost. At t3, according to evidence available to X, the 50000 can be saved at the cost of a further 10000. And the same thing happens again and again, with continued losses, and without X getting any closer to saving the 50000. Some favour *Quota* because it avoids the implication that iterated permissible mistakes like this are possible.⁷

This argument for *Quota* also fails. First, it is hard to see why the possibility of iteration has implications in circumstances where it does not occur. We have already

⁷ See, for example, Moellendorf 'Two Doctrines of *Jus ex Bello*', 664-6. Seth Lazar discusses the idea at greater length in 'Moral Sunk Costs', unpublished ms.

noted that acts can be fact-relative wrong or evidence-relative wrong. First begin with the question whether the possibility of iteration makes X's act at t2 fact-relative wrong. Suppose that X continues to fight in *Early Losses 2* at t2. It kills 10000 and saves the 50000. It is difficult to see how the possibility of iteration can make doing this fact-relative wrong where iteration does not occur.

But if the possibility of iteration cannot make X's conduct fact-relative wrong, it is also difficult to see why it should make it evidence-relative wrong. X had evidence at t2 that there was some possibility that the 50000 would not be saved without many iterations of the same set of acts under similar evidential circumstances - there is always some chance that this is so. But this is just one piece of evidence that X has at each successive stage of the war to determine whether continuing is evidence-relative permissible. It must be weighed against other possibilities, such as the possibility that X will save many more people than expected, or kill far fewer people than expected, if it continues.

Perhaps it might be argued that the possibility of iteration is especially important, because if it is permissible to iterate one's decision the war can cause catastrophic harm. But the probability that this will occur may be tiny, and friends of *Quota* cannot think that a small probability of catastrophic harm is necessarily decisive. There is always some chance that a small war will trigger a catastrophic event, such as a nuclear war, without the aim of the small war being achieved, yet friends of *Quota* are not pacifists. And, of course, there is always some probability that catastrophic harm will occur if one does not go to war, or does not continue a war that one has started.

It is also not clear why the potential for catastrophic losses through iteration supports *Quota* in particular. There is an epistemic risk that 10000 will be killed without X making progress to save the 50000 at each successive stage in an iterated series. Why should the possibility of catastrophic losses demand that X stops at t2 rather than at t1, t3, or some point further down the line?

A second response draws on the idea that the problem of iteration is not restricted to single wars. It can arise across decisions to start different wars, and decisions not to start different wars. Because of this, if there is a problem of iteration,

it arises for *Quota* as well as *Prospect*. Indeed, it arises for any evidence-relative principle in the ethics of harm. This should incline us to think that there is no problem of iteration.

Quota requires a country not to pursue the particular war that it has started when it causes a certain number of deaths, but it has no implications for decisions to start new wars. Now suppose that *Quota* is true, and X thus stops fighting in *Early Losses*. A further 50000 people are then threatened in another country, Y2. X goes to war, suffers early losses in battles resulting in the deaths of 10000 people, and then stops fighting. A further 50000 people are then threatened in a further country, Y3, and so on. X can comply with *Quota*, and yet X will kill a very large number of people for no benefit at all. Those who have the intuition that X ought to stop in the intra-war case will surely have the same intuition in the inter-war case.

Perhaps some will claim that this is enough to show that *Prospect* is false, though. It is just that past harms caused should be discounted in both the intra-war and the inter-war case. But in response, we can extend the objection further. Just as there is no reason to think that the problem of iteration is specific to inter-war cases, there is no reason to think that it is specific to causing harm over allowing harm. There can also be a problem with successive decisions not to go to war. In such cases, the decision not to go to war may result in successive failures to save lives where, it turns out, these lives could have been saved at no cost.

Suppose that X starts five wars, suffers early losses in each, and so kills 50000 without achieving anything. X then stops fighting as friends of *Quota* now recommend. Y6 then threatens to kill 50000 people. X expects to be able to save the 50000 at the cost of 10000 lives. X does nothing, the 50000 are killed, and afterwards it becomes clear that X could have saved the 50000 without killing anyone – if they had crossed the border into Y6, official documents reveal, Y6 would immediately have surrendered. Y7 then threatens to kill 50000 people. Again, X expects to be able to save the 50000 at the cost of 10000. Again X does nothing, the 50000 are killed, and it again becomes clear that the 50000 could have been saved without killing anyone. And so on. The intuition that X should, at some point, go to war is just as strong as the intuition that X should stop in the intra-war and the inter-war cases.

This shows that if there is a problem of iteration, it is a general problem in evidence-relative morality. For almost any putative evidence-principle in the ethics of harm, complying with the principle can have unexpected very bad consequences, or the expected bad consequences can occur without anything good occurring that would have justified them. So complying with the principle over and over again can have very bad consequences over and over again, with extremely bad consequences overall, and without any good occurring that would have justified them.

Perhaps it might be argued that the problem of iteration should lead us to favour a principle that does not permit iteration of the same mistake. That would count against *Prospect*, because *Prospect* permits X to do the same thing that it has already done. But we have no reason to favour a principle that allows many different fact-relative mistakes to be made in succession over one that allows the same fact-relative mistake to be made in succession! And this is especially so if abiding by the principle that recommends altering our conduct increases the probability that our successive acts will be gravely wrong in the fact-relative sense.

It is hard to grasp iteration cases, given that the bad consequences of our previous evidence-relative permissible acts often give us new evidence that alters what it is evidence-relative permissible to do. This makes our intuitions about such cases untrustworthy. Once we see that the possibility of multiple mistakes arises for any evidence-relative principle, we should conclude that any inclination that we have to revise our evidence-relative principles in the light of the possibility of iterated mistakes is fully explained on epistemic grounds. Thus *Iteration* fails to support *Quota* (or, for that matter, *Discount*).

II. The Problem of Small Costs

So far we have seen that three arguments that have been offered for *Quota* fail to support it. I now offer this decisive argument against it: *Quota* implies that it is wrong for X to achieve a very great good by causing a very much smaller amount of harm at t_2 in cases where this is clearly false. For example, *Quota* sometimes has the unpalatable implication that it would be wrong to save the 50000 at the cost of a

single extra life, because too much harm has already been caused in the attempt to save the 50000.

To illustrate the objection, suppose that there is some number between 10000 and 12000 that is the precise tipping point that makes saving 50000 disproportionate. Early deaths caused during the war put X just below the tipping point. X can now save the 50000 in a way that will kill one extra person. *Quota* implausibly implies that doing this is wrong.

Moellendorf responds that the fact that the whole good can now be achieved at such a small cost demonstrates that the proportionality judgement at t_1 was mistaken. So he rejects a premise of the argument.⁸ This response cannot be right. If there is a precise tipping point, causing n deaths would be proportionate and causing $n+1$ deaths would be disproportionate. For any value of n , it is possible that n deaths are caused in early battles achieving nothing, but that the whole good can be achieved at t_1 at the cost of an extra life. As there must be some value of n that is the proportionality threshold, and the problem of small costs can arise whatever that value is, the fact that the problem of small costs arises does not show that there is a mistake about the initial proportionality judgement.

Friends of *Quota* may respond that the problem of small costs does not arise because there is indeterminacy at the proportionality threshold.⁹ About *Early Losses*, they may claim that there are numbers between 10000 and 12000 where it is indeterminate whether the war is proportionate. This does not adequately answer the objection either. It is hard to believe that the judgement whether *Quota*, *Discount*, *Prospect*, or *Addition* is true depends on the extent of indeterminacy at the proportionality threshold. The problem of small future costs brightly illuminates the fact that past deaths that one has caused cannot make it wrong to achieve a very great good at very little cost. The possibility that the proportionality threshold is indeterminate is not responsive to this objection.

⁸ See Moellendorf 'Two Doctrines of *Jus ex Bello*', 667-8.

⁹ Seth Lazar suggests this possibility in 'Moral Sunk Costs'. He attributes the point to discussion with Emily McTernan.

Furthermore, this response only makes the implications of *Quota* more plausible where X's acts at t1 put it near the proportionality threshold. The response fails in cases where X's earlier conduct clearly puts it above that threshold, but the objection seems just as forceful in that case.

Consider:

Early Big Losses: As *Early Losses* except that 75000 are unexpectedly killed before t2. The 50000 can now be saved at an expected cost of one life.

Quota implies that it would be wrong for X to continue to fight. Here is why. It would clearly be disproportionate for X to go to war at t1 to save 50000 at the cost of 75000 lives. It would be even more powerfully disproportionate at t1 to save 50000 at the cost of 75001 lives. We are not permitted to kill many more people than we save, other things equal, on any sensible view. There is thus no question of revising the initial proportionality calculation to explain why saving the 50000 seems permissible. And even if there is indeterminacy at the proportionality threshold, this war as a whole does not fall in the range where this is true. Therefore, *Quota* implies that it is wrong for X to save 50000 lives at the cost of one life at t2. This implication is so implausible that *Quota* should be rejected simply on this basis.

One way to reinforce the conclusion that this is implausible draws on an idea explored in the previous section. In determining what X should do, we should consider X's acts together if it either continues or stops. If X continues it will have killed 75001 people to save 50000. This is obviously fact-relative wrong, and gravely so. And at t2 it has decisive evidence that it will have committed this serious fact-relative wrong if it continues. But suppose that it stops. It will then have killed 75000 people to save no one. And killing 75000 people to save no one seems even more gravely wrong than killing 75001 people to save 50000.¹⁰

¹⁰ Fabre acknowledges that reasoning of this kind seems to have force in *Cosmopolitan Peace*, 40.

Note that the following options were available at t1 (though X did not know it):

- 1) Kill no one, resulting in the death of 50000
- 2) Kill 75001, but save 50000
- 3) Kill 75000, but save no one.

Clearly, X ought to have picked 1) if it were to have known that these were its options. At t1, all three options were available to X, and that explains the intuition that X acts wrongly, in a sense, by killing 75001 to save 50000. This intuition may draw people to *Quota* – it explains why X's acts together are fact-relative wrong. But X cannot now make it true that its conduct as a whole kills no one, resulting in the deaths of 50000. Nothing like 1) is an option for X at t2 in *Early Big Losses*. It can make it true that it has killed 75001, saving 50000 or it can make it true that it has killed 75000, saving no one. These are similar to options 2) and 3) that X had at t1. Between 2) and 3) at t1, 3) is much worse than 2). Similarly, at t2, X ought to make it true that it has killed 75001 saving 50000, rather than making it true that it has killed 75000 saving no one. By continuing to fight X mitigates the gravity of its overall wrongdoing when compared with stopping.¹¹

Thinking in this way also helps to support *Prospect*. Here is a natural way to select between 2) and 3). One thing that is equal between 2) and 3) is that 75000 die as a result of X's war. And as these deaths have already occurred at t2, we should assume that the very same 75000 people will be killed. It seems natural, then, to treat these deaths as providing an equally powerful reason against 2) and 3). If this is right, we should chalk these deaths off for the purposes of deciding what to do, and consider only the remainder. We should take a similar approach to X's acts at t2. It cannot affect the lives of those it has killed, so we should chalk these deaths off for

¹¹ This is a way of illuminating more brightly a similar idea in D Rodin 'The War Trap: Dilemmas of *jus terminatio*' (2015) 125 *Ethics* 674, 686-92.

the purposes of deciding what to do, and consider only the remainder. And that is what *Prospect* recommends.

III. How Do Previous Deaths Count?

We have seen that past deaths cannot have the same weight as future deaths in the decision whether to continue a war. Friends of *Discount* might respond that *Prospect* is implausible, though, because it implies that past deaths do not count at all. For example, Fabre, in considering the implications of *Prospect* for a case that is similar to *Early Losses* writes:

Let us...consider the fate of the ten thousand agents who die between t1 and t2. At t1, those deaths are regarded as a bad to be weighed relative to the good the war would bring if A initiated it and thus count as a reason against going to war – albeit not a sufficiently decisive one. At t2, those deaths simply do not count as a bad any more to be weighted relative to the good the war would bring if A continued it. The worry is that, on this view, proportionality would lose most of its bite as a constraint against killing.¹²

Fabre goes on to suggest that *Prospect* implies that proportionality does not appropriately constrain war.

Discount might then seem preferable because it gives some weight to past losses, and so does more to constrain war. But there two are decisive objections to the idea that past deaths count just like future deaths, though with less weight, in determining whether it is proportionate to continue fighting.

First, the objection concerning small future costs can be extended to show that a version of *Discount* that makes this assumption about relative weighting is false. To see this, consider:

¹² See C Fabre 'War Exit' 637.

Early Catastrophic Losses: as *Early Big Losses*, except that the early battles have catastrophic results because nuclear weapons are unexpectedly used, killing several million people. X can still save the 50000 at the expected cost of one life.

The version of *Discount* under consideration implies that X ought not to save 50000 at the cost of one life in *Early Catastrophic Losses*, because each of the several million deaths counts against continuing to fight, and the number of deaths is so large that it will outweigh the value of the 50000 lives that will be saved. Versions of *Discount* that have this implication are clearly false.

Second, past deaths cannot possibly count against continuing to fight in the same way as future deaths. Here is why. The deaths that X will cause can, on their own, make it wrong for X to save the 50000. If it will kill too many, it ought not to save the 50000. The deaths that X has already caused cannot play this role. If it can save the lives of 50000 people without harming anyone at all, it is obviously permissible to do this, no matter how many deaths it has already caused.

This suggests that if past deaths count against continuing to fight, they can only do so indirectly. They might do so either by strengthening the reasons against killing further people or by weakening the reasons to save the 50000. But they cannot provide reasons against continuing to fight on their own.

IV. Intrapersonal and Interpersonal Morality

Here is an argument that past harms strengthen the reasons against killing more people to save the 50000. It draws on the relationship between intrapersonal and interpersonal cases.¹³

¹³ Fabre and I discussed the relationship between intrapersonal and interpersonal cases in relation to sunk costs on PEA Soup. Following that discussion, Fabre considers the relationship between intrapersonal and interpersonal cases in *Cosmopolitan Peace* 39-40. I am not completely confident of her current view. On the

The fact that in the past X has harmed an innocent bystander, A, in the course of attempting to save Y plausibly makes a difference to the permissibility of inflicting more harm on A in a further attempt to save Y. This may be true because it matters that A was harmed by X, or because A has already been harmed whether or not that is by X, or even simply because A is badly off. This is an intrapersonal case: it concerns the implications of the harm that a person has suffered for the permissibility of harming that same person more. Defenders of *Discount* might argue that although there is a difference between intrapersonal and interpersonal harm aggregation, this is a difference of degree, not kind. Therefore, the fact that X has harmed A in the course of attempting to save Y in the past makes a difference to the permissibility of inflicting harm on another innocent bystander, B, in a further attempt to save Y.

Intrapersonal cases may have some implications in war – those who suffer as a result of the deaths of some people may suffer more from the deaths of others. In this way, past deaths can count against the decision to keep fighting. Furthermore, some people who are injured in war may be vulnerable to be harmed further if war continues, and again that may make a difference to the permissibility of continuing to fight. I set these issues aside.

Our intuitions in intrapersonal cases do seem more powerful than our intuitions in interpersonal cases. However, there are explanations of our intuitions in intrapersonal cases that do not extend to interpersonal cases, and for this reason friends of *Discount* are wrong to rely on them. For example, consider the prioritarian view that there is a stronger reason to benefit a person, or not to harm her, if our conduct will leave her worse off than if it will leave her better off.¹⁴ On this view, it is

one hand she seems to offer the intrapersonal case to defend her rejection of something like *Prospect*. But then she acknowledges that it does not have force in pure interpersonal cases, and goes on to suggest that intrapersonal cases are common in war.

¹⁴ In Derek Parfit's classic statement of it, prioritarianism is the view that benefiting people matters more the worse off these people are. We are concerned with harming

not only the size of the difference that our conduct makes to a person that determines what we should do; it is the absolute level of welfare that a person will have if we act or refrain from acting. This idea helps to explain our intuitions in intrapersonal cases, but has no implications in interpersonal cases. The fact that one person has been left very badly off as a result of our actions has no implications for what we should do to others as far as prioritarianism is concerned.

Of course, the prioritarian argument, if it is successful, implies that it is harder to justify harming people who are badly off no matter why they are badly off – whether that is because of past harms that have been inflicted on them or not – and no matter the identity of the person who will harm them. But then, it is not at all obvious that it does make a difference to the permissibility of harming a person that they have suffered harms, or suffered harms by the person who might now harm them again.

A similar thing is true of a further argument that applies in intrapersonal cases – one concerning personal prerogatives. A person is entitled to use her personal resources to promote her own ends, within certain limits. How the limits of personal prerogatives are determined is a difficult question that I cannot address here. But it is plausible that the costs that a person is required to bear for the sake of some goal depend in part either on whether the person has been harmed for the sake of that goal, or harmed independently of that goal, or is badly off. Furthermore, it is plausible that the extent to which it is permissible to harm a person for the sake of some goal depends in part on whether she would be required to bear costs for the sake of that goal. Again, this argument applies in the intrapersonal case, but it has no implications in the interpersonal case.

Of course, some might deny that these ideas fully explain our intuitions in intrapersonal cases. But they are sufficient to undermine the argument that *Discount*

rather than with benefiting. See D Parfit 'Equality and Priority' in M Clayton and A Williams *The Ideal of Equality* (London: MacMillan, 2000). Prioritarians, though, will surely also believe that our reasons against harming people are stronger the worse off these people will be.

gains support from intrapersonal cases. The more powerful intuitions that we have in intrapersonal cases can provide no support for a view like *Discount* if there are plausible explanations of those intuitions that do not apply in interpersonal cases. And there are. Thus we do better to consider interpersonal cases directly.

V. Do Past Harms have Indirect Significance?

We have seen that if past deaths are morally significant for future conduct in *Early Losses*, they are only indirectly significant. The two most natural indirect ideas are that past deaths weaken the rights of the 50000 to be saved, or that they strengthen X's duty not to kill more people.

i) *Are the Rights of the 50000 weakened?*

One idea is that the right of the 50000 to be saved is weakened by the fact that harm has been caused in the course of an attempt to save them. Some might draw on the idea that after a certain number of deaths are caused in the course of attempts to rescue them, enough has been done for them.

A radical, and implausible, version of this idea is that the rights of the 50000 can be exhausted by deaths caused in the course of attempts to save them. This view implies that if enough people are killed in the attempt to save the 50000, they have no right to be saved even if they can be saved at no cost to anyone.

To reinforce the conclusion that this view is implausible, notice that it implies that the 50000 have lost their right to be saved from being wrongfully killed by Y's officials because of a terrible accident that occurred in the course of an attempt to save them. They have done nothing to lose their rights. The idea that their rights to be saved from the lethal threat they face is completely vitiated simply by their bad luck at others trying and failing to save them in a way that causes harm to others is hard to accept.

This leaves open the possibility that their rights to be saved are weakened but can never be negated by the harm that has been caused for their sake. And if that is

true, it might be argued that they can no longer demand that attempts are made to save them where these attempts would risk the lives of many innocent people, even if they would have had such a demand were past harms not to have been caused for their sake.

Even this weaker idea seems false. I suspect that any intuitive force in the idea that enough has been done for the 50000 draws on the attractive and more general idea that there are limits to what we have to do for others. But I think that a proper understanding of the latter idea does not imply that the rights of the 50000 are weakened.

Earlier, I offered two explanations of the wrongness of harming a person who has already been harmed: prioritarianism and agent-relative prerogatives. These ideas also help to explain the general idea that there are limits to the burdens we need to bear for the sake of others. But we have already seen that neither of these ideas support *Discount*.

Discount relies on some further version of the idea that there are limits to how much we must do for others. But it is not clear what this idea is. Suppose that one person tries to rescue me and fails. That person might plausibly claim that she has done enough for me if the cost is large enough. Those who are harmed as a side-effect of the failed attempt might also plausibly claim that they cannot be expected to bear any further cost for my sake.

But why should this have any bearing on what I am owed by those who have not yet tried to help me, or those who have not yet been harmed in any attempt to save me? My circumstances have not changed as a result of the attempt: I am in just as much peril as I was before. These other people have not been affected either. So it is hard to see how the failed attempt to rescue me has any bearing on the strength of my right to be saved.

ii) *Are X's Duties not to Harm Others Strengthened?*

Another possibility is that X's duty not to harm others is strengthened by the deaths that it causes. A broad version of this idea is that X's general duty not to harm others

is strengthened, whether or not harm is inflicted for Y's sake. The more deaths there are on X's ledger, it might be argued, the more stringent is its duty not to kill in the future.

This general view comes in several flavours. The broadest view counts all deaths on X's ledger, whether or not inflicting those deaths was fact-relative wrong. This view is especially implausible. Consider a country that engages in very many just wars, achieving a great deal of good. The harm inflicted in each war, considered on its own, would have been proportionate. The broad view implies that each successive war becomes harder to justify, making some of the later wars disproportionate. There is little appeal in the idea that a country that is especially good at fighting proportionate wars faces increasingly high proportionality hurdles.

A slightly narrower view counts only those deaths that it was fact-relative wrong for X to cause. The more such deaths that X causes, the more stringent X's duty is not to cause further deaths. But it is also hard to support this view. It is only permissible to cause a death if in doing so one will achieve some good aim (or has good prospects of doing so). The aim under consideration in cases such as *Early Losses* is the prevention of wrongful death by others. The view under consideration implies that those who wrongly cause deaths (in the fact-relative sense) should be more inclined to allow such future deaths to occur than to cause such future deaths. It is hard to see why this should be so: why should past fact-relative wrongdoing make a difference to the stringency of the obligation not to cause more deaths, but not to the stringency of the obligation not to allow more deaths?

A still narrower view counts only those deaths that it is fact-relative wrongful to cause in this particular war. This view is also problematic. There is a difficult metaphysical question how to individuate wars. It is often hard to know whether some fighting counts as the continuation of one war, or as the beginning of a new war. Fortunately, those working on the morality of war don't seem to need to answer that question, as it seems that there is nothing morally significant about how it is answered. The view under consideration, though, would make this seemingly irrelevant question morally significant.

Finally, we might consider the view that the deaths that matter are those that are caused in the service of the end of saving the 50000. But this view is very close to *Weakened Rights* that we considered in the previous subsection, and in the light of the arguments offered there it is hard to see how to support it.

Overall, it proves very difficult even to state a version of *Discount* that has plausible implications. I have explored several arguments to support it, but none of them seem convincing. And we can show why its apparent appeal is illusory. Overall, I think we should reject it.

VI. Doubts about Addition

This leaves us with *Prospect* and *Addition*. Recall that *Addition* is the view that deaths to innocent people caused in the past make it easier to justify harming others, so that killing more innocent people in the future is permitted to achieve the same goal than would have been the case had these past deaths not occurred.

Arguments for *Addition* are even harder to find than arguments for *Discount*. Like *Discount* in order to be at all plausible *Addition* needs to be carefully interpreted. If it is not, it will suffer from the opposite problem to *Discount* that we considered in Section III: if the death toll is high enough, the reasons against killing more people for the sake of a goal will become weaker and weaker, so that it will become permissible to kill very many people in order to save far fewer people. This suggests that if lives lost make it easier to justify killing, they only do so somewhere near the proportionality threshold. Let us consider a version of *Addition* of that kind.

i) *Disincentives*

One argument that might be offered in favour of *Addition* is that those with the attitudes that *Addition* recommends them to have incentivise their opponents to minimize casualties in war.¹⁵

To see why, recall *Early Losses*. Whether X's war is, as a whole, proportionate, depends on whether more than 10000 are killed to save the 50000. Some of the deaths that result from X going to war are caused directly by X, some directly by Y, and some directly by third parties. X causes the deaths of the latter two groups indirectly – X's decision to go to war affects the acts of Y and others, resulting in some deaths. There is a difficult question about whether all of these deaths count equally in the proportionality calculation, given the role of intervening agents. For the sake of simplicity, let us assume that they do.¹⁶

At t1, the deaths that Y will cause if X continues to fight help to make continuing to fight disproportionate. Now suppose that X has the attitude that it will treat deaths caused as reasons against continuing to fight, as *Quota* recommends. Y then has an incentive to kill, for the more people it kills in response to X's decision to go to war, the more likely it is that X's war will be disproportionate. If X respects *Quota*, the sooner Y makes it true that X has reached its quota, the sooner X will stop fighting. This may seem a good reason to reject *Quota*.

If it is, it is also favours accepting *Addition*. If X has the attitude that it will treat deaths caused as reasons to continue fighting, as *Addition* recommends, Y has an even stronger disincentive to kill. For the more innocent people that Y kills in the course of the war, the more X will be inclined to keep fighting.

I suspect, though, that this argument does not support *Addition*. It is an argument that X should have the attitude of treating *Addition* as true rather than an argument that *Addition* is true. The reason X should have the attitude of treating *Addition* as true is that its having this attitude maximises its chances of saving the

¹⁵ Here I draw on T Kelly 'Sunk Costs, Rationality, and Acting for the Sake of the Past' (2004) 38 *Noûs* 60, 65-70.

¹⁶ For arguments that what is assumed is true, see V Tadros 'Permissibility in a World of Wrongdoing' (2016) 44 *Philosophy and Public Affairs* 101.

50000 at as little cost as possible. But the fact that its having this attitude would have this effect cannot alter the proportionality calculation as such. Rather, X has a reason to have an attitude of treating past deaths as reason-conferring even though they are not, because doing so minimizes losses, and thus makes X more likely to comply with its actual moral obligations.¹⁷

ii) *Redemption*

Earlier, I explored a way of supporting *Discount* by drawing on intrapersonal cases. Some may be tempted to do the same in defending *Addition*. In intrapersonal cases, many people are inclined to honour sunk costs. If we incur some cost in pursuit of a plan and then abandon the plan, the cost will have gone to waste. Whilst many people think that the fact that many people do this is a sign of our irrationality, others disagree. They think that there is value in redeeming the costs that we have borne for the sake of a certain goal; something that we do by achieving the goal.¹⁸ Again, if the difference between intrapersonal and interpersonal cases is only one of degree some support might be found for *Addition*.

Furthermore, the intrapersonal argument in favour of *Addition* might seem better than the intrapersonal argument in favour of *Discount*, at least in one way, because those who think that we have reasons to honour sunk costs explain why bearing a cost for the sake of some end matters as such, rather than instrumentally.

Some think that the fact that a person has sacrificed something for the sake of some end is a reason for that person or others to pursue that end. The reason is that

¹⁷ For a similar way of understanding the relationship between our reasons to act and our reasons to have certain attitudes that determine how we will act, compare D Parfit *Reasons and Persons* (Oxford: OUP, 1984) s.5.

¹⁸ For subtly different views, see S Keller 'Welfare and the Achievement of Goals' (2004) 121 *Philosophical Studies* 27; T Kelly 'Sunk Costs, Rationality, and Acting for the Sake of the Past'; D Portmore 'Welfare, Achievement, and Self-Sacrifice' (2006-8) 2 *Journal of Ethics and Social Philosophy* 1.

their sacrifices will then not have been in vain.¹⁹ In the central cases, one person, A, makes a sacrifice for the sake of some goal, *g*, but fails to achieve *g*. Either A, or a second person, B, can add to A's contribution, and *g* will be achieved. If *g* is achieved, A's earlier sacrifices will not have been in vain, and this gives A a reason to achieve *g*, and it gives B such a reason too, especially if A cannot achieve *g*. This is part of a wider set of cases where one has reason to act in a certain way because that will bestow meaning or value on one's previous actions or efforts.²⁰

Jeff McMahan plausibly argues that deaths to combatants on the just side of a war can make achieving a just cause more valuable because achieving it would redeem the sacrifices of those who have died in the course of attempting to secure that cause.²¹ Some might then argue that because achieving the just cause would have greater value, greater costs can be inflicted on innocent people as a side-effect of its pursuit. Although McMahan does not think that wide-proportionality is affected by redemption, he admits that he lacks a fully satisfactory explanation why; he describes his view – that redemption can make a difference to narrow proportionality, but not to wide proportionality – as paradoxical.²²

Even if there are reasons of redemption, they cannot contribute to a defence of *Addition*, for *Addition* is not concerned with the deaths of those who are pursuing just aims, but only the deaths of those killed as a side-effect of the pursuit of those aims.

¹⁹ Abraham Lincoln relied on this idea in the Gettysburg Address when he said: 'It is for us the living, rather, to be dedicated to the great task remaining before us – that from these honoured dead we take increased devotion to that cause for which they gave the last full measure of devotion – that we here highly resolve that these dead shall not have died in vain.' I am grateful to an associate editor of *Philosophy and Public Affairs* for the quote.

²⁰ For good discussion, see J McMahan *The Ethics of Killing: Problems at the Margins of Life* (Oxford: OUP, 2002) 174-85; Kelly 'Sunk Costs, Rationality, and Acting for the Sake of the Past'.

²¹ 'Proportionality and Time' 710-17.

²² 'Proportionality and Time' 718.

As we will see in a moment, it is hard to extend the idea to support *Addition*. But I also think that McMahan's view that wide-proportionality is more generally unaffected by reasons of redemption can be defended. Here is how. Fully just combatants intend both to secure the just cause, but also to respect the value of the lives of those they will harm as a side-effect of pursuing that cause. Those aiming to redeem their losses must respond appropriately to both of these attitudes when deciding how to redeem their sacrifices. But if they kill more people than would otherwise have been permissible to secure the ends of those who have been killed they will fail to do so.

To illustrate, suppose that Phase 1, a group of *X*'s well-motivated combatants, aims to contribute to the saving of the 50000 and they are killed in the process. Phase 2 can then save the 50000, but they will kill 12000. I have stipulated that it would have been disproportionate for Phase 1 to save 50000 at the cost of 12000. This is because of the value of the lives of the 12000. As Phase 1 would not have saved the 50000 at this cost, it is hard to see how saving the 50000 at that cost redeems the sacrifices of members of Phase 1. True, they were pursuing the aim of saving the 50000. But they aimed to do so only in a way that would show proper respect for the lives of those that they might otherwise kill in the process.

Given that the 12000 have done nothing to alter their moral status, it is hard to see how members of Phase 2 would respect the ambitions of members of Phase 1 by saving the 50000 at the cost of 12000 lives. They would disrespect those who have died by treating the lives of the 12000 as less significant than the 50000 in a way that those in Phase 1 would have rejected in their own plans to save the 50000. Obviously, the redemption argument cannot be rescued by pointing to the fact that Phase 1 would have killed disproportionately to save the 50000 – we don't have reasons to redeem wrongful plans. So it seems that the redemption argument fails, even when we focus on combatants who die in pursuit of the just cause.

How, then, can redemption make a difference to narrow proportionality? It might seem that it cannot, for just combatants also have the ambition of abiding by narrow proportionality considerations. One answer is that deaths earlier in the war can alter the grounds of liability of those fighting unjustly at a later time. The

achievement of the just cause would help to redeem the losses suffered by combatants who died earlier in the war. This gives combatants fighting on the unjust side an extra reason to desist – not only does their fighting threaten to prevent the initial just cause from being realized; it threatens to prevent the losses suffered by those who have died on the just side from being redeemed. Their liability to be harmed, then, may be grounded not only in their responsibility for a threat to the achievement of the initial just cause but also in their responsibility for a threat to the achievement of redemption. I am not sure even whether this view is true, but it is plausible.

As McMahan notes, though, even if this argument is wrong, and redemption makes a difference to wide as well as narrow proportionality, it is difficult to see how it can be extended to cases like *Early Losses*. The value of redemption seems most plausibly explained by the respect that we have for the decisions that rational agents make to pursue certain ends and for the sacrifices they make in pursuit of those ends.²³ We can make the projects that they have sacrificed a great deal for successful by completing those projects ourselves. Some also argue that this idea can improve the welfare of those who have made sacrifices in pursuit of the relevant goals. This is plausible because it is plausible that our welfare depends in part on our achievements.²⁴

This idea applies more naturally to combatants than to innocent people killed as a side-effect in war. Although their lives have been lost in the course of an attempt to save the 50000, they were not pursuing the goal of saving the 50000. So the idea that their deaths will have been sacrificed in vain if X does not continue to fight has

²³ See, also, Kelly ‘Sunk Costs, Rationality, and Acting for the Sake of the Past’, 78. For the view that it is sacrifices made, rather than effort expended, that ought to be redeemed see Portmore ‘Welfare, Achievement, and Self-Sacrifice’.

²⁴ See Portmore ‘Welfare, Achievement, and Self-Sacrifice’; Keller ‘Welfare and the Achievement of Goals’. For the view that achievements contribute to welfare more generally, see T M Scanlon *What We Owe to Each Other* (Cambridge, Mass.: Harvard UP, 1998) ch.3.

much less force in their case; even if it has some force, I doubt that it has enough to make a difference in matters of life and death.²⁵

Conclusion

Quota and *Discount* have a certain initial appeal, but that appeal is largely illusory. That does not show that these views are wrong. *Quota* is clearly wrong, though. It is more difficult to decide whether *Discount* is wrong. A defence relies on some version of *Weakened Rights* or *Strengthened Duties*. But it is difficult to see how to defend these views. The best arguments for *Addition* also seem to fail. Overall, I conclude that we have most reason to accept *Prospect*.

²⁵ See, also, McMahan 'Proportionality and Time' 713.