

Personal and Omnipersonal Duties

Derek Parfit

ABSTRACT: This paper's main aim is to discuss the relations between our duties and moral aims at different times, and between different people's moral aims and duties. The paper is unfinished because it was written as part of an intended chapter in the third volume of my book *On What Matters*, and I later decided to drop this chapter. That is why this paper asks some questions which it doesn't answer. But though this paper does not end with some general conclusions, it defends some particular conclusions.

When Ross sums up his main objections to Act Consequentialism, he writes that any such theory “ignores, or at least does not do full justice to, the highly *personal* character of duty.”

There are at least two senses in which some duty might be personal. Some duties are personal, Ross claims, in the sense that different people have duties to try to achieve different aims. Some of these duties may also be personal in the sense that we have these duties only to certain other particular people. Of duties that are personal in both these senses, one example is our duty to keep promises. We should each have the aim of keeping our own promises, and we have this duty only to those particular people to whom we have made these promises. We have similar, doubly personal duties to care for own children or other close relatives, to show gratitude to those who have benefited us, and to make reparations to those people whom we have injured.

Some other duties are not personal in the second of these ways, because we have these duties to everyone. We have such duties not to kill or injure people, and not to treat people in certain other ways, such as deceiving or coercing them. When some duty is not personal in this second sense, because we have this duty to everyone, this duty might either be personal in the first sense, by giving different people different aims, or be *omnipersonal*, in the sense that it gives everyone the same common aim. We all have a duty, for example,

Derek Parfit is an Emeritus Fellow at All Souls College in Oxford. He specializes in topics in ethics and personal identity, and is widely considered to be one of the most influential modern moral philosophers. He is a Fellow of both the British Academy and the American Academy of Arts and Sciences and a recipient of the Rolf Schock Prize. Parfit's first book, Reasons and Persons, famously introduced the Repugnant Conclusion in population ethics, and his most recent book, On What Matters, follows up his work on values and normativity. E-mail: derek.parfit@all-souls.ox.ac.uk.

to save the lives of strangers, when the cost to us would be below some rough threshold. This duty would be personal in the first sense if each of us ought to have the aim that we ourselves will save as many lives as possible. This duty would not be personal in either sense if everyone ought to have the common aim that as many lives as possible are saved.

Of the duties that are personal in the sense that they give different people different aims, some are also temporal, or *time-relative*, in the sense that they give people different aims at different times. Doctors and lawyers, for example, may have such time-relative duties to those who are now their patients or clients. Other personal duties are temporally neutral, by giving people the same aims at all times.

In asking which duties are of these kinds, we can first compare two versions of what we can call *the Personal Life-Saving Principle*. According to

the Time-Relative Principle: We have a prima facie duty to act in the way that, of our possible present acts, would save the most people.

According to

the Temporally Neutral Principle: We have a prima facie duty to act in the ways that, in our life as a whole, would save the most people.

Suppose that, in

Case One, you could either save one person now, or save two people later.

There may be some versions of this case in which we could plausibly believe that you ought to save one person now, even though your act would make it impossible for you to save two people later. If you are a doctor, for example, you might have a duty to save one of your present patients, even if you would have to use some scarce drug which, if kept, would enable you to save two of your future patients. We may similarly believe that we ought to give priority to rescuing certain miners who are now trapped underground, though we know that the money that would be spent on this rescue operation would later save more lives if this money was spent instead on preventive safety measures.

If we have these beliefs, however, we are not assuming that a mere difference in timing makes a moral difference, so that we ought to save one person now rather than saving two people later. We are appealing to what we believe to be some morally relevant ways in which we might be more closely related to the people whose lives we could now save. To choose between these versions of the Life-Saving Principle, we should consider cases in which we have no such special relations to the people whose lives we could save. Suppose that you are the only person with the knowledge that would be needed to save three people's lives. One of these people is nearby, but the other two are far away. If you spent the time that would be needed to save the first person's life, you would be unable to reach the other two people in time to save their lives. Most of us would believe that, since you are not more closely related to the person whose life you could save now, the fact that you would be saving this person now does not give him or her a moral claim that is as strong as the combined claims of the two people whom you could later save. Such mere differences

in *when* we could save some people's lives do not affect the strength of people's moral claims to have their lives saved by us. In this example, as the Temporally Neutral Principle implies, you ought to save the other two people's lives.

We can next compare two versions of *the Personal Harm Principle*. According to

the Time-Relative Harm Principle: We have a prima facie duty to act in the way that, of our possible present acts, would kill or harm the fewest people.

According to

the Temporally Neutral Harm Principle: We have a prima facie duty to act in the ways that, in our life as a whole, would kill or harm the fewest people.

Suppose that, in

Case Two, you have earlier acted wrongly, in some way that threatens to kill five people. You could save these five people's lives, but only by acting in some other way whose foreseen side effect would be to kill some other person.

On the Temporally Neutral Principle, as stated above, you would have a prima facie duty to do what would save these five people's lives, since you would thereby make it true that, in your life as a whole, you would kill four fewer people. On the Time-Relative Principle, it would be wrong for you to do what would kill this one person, since you could now act in a way that killed no one. You ought not to kill this person now even though you would thereby make it true that, in your life as a whole, you would kill fewer people.

Though this disagreement raises several interesting questions, I shall not discuss these questions here. It is easier to try to decide first what we ought to do in cases in which no one acts wrongly, turning only later to questions about what we ought to do in cases in which some people act wrongly. Except when I say otherwise, the principles that I discuss should all be taken to apply only to cases in which no one acts wrongly.

In one such example,

Case Three, some nuclear power plant has been damaged by a meteorite, so that some reactor is in danger of exploding, in a way that would kill about a million people. You are an engineer in this power plant, and are the only person who could prevent this explosion, thereby saving these million people. You could either prevent this explosion now, in a way that would kill one other person, or prevent this explosion later, in a way that would kill two other people.

On the Time-Relative Harm Principle, it would be wrong for you to prevent the explosion now, since this act would kill one person, and you don't need to act in this way to prevent

the explosion. You know that, if you don't prevent this explosion now, it will later become your duty to prevent this explosion in a way that would kill two other people. But this fact could not justify your killing someone now. On the Time-Relative Principle, your duty at any time is to kill as few people as possible at this time.

These claims are clearly false. As the Temporally Neutral Principle implies, you ought to do what would kill one person now, since that is your only way to prevent its becoming your duty to kill two people later. You should have the temporally neutral aim of doing what, in your life as a whole, would kill the fewest people.

I am not assuming here that, in cases in which there is no wrong-doing, we would all have a duty to do what would kill the fewest people. There may be some people who have stronger moral claims not to be killed by us. But no such claim applies to *Case Three*, which involves a mere difference in timing.

We can next revise this example so that it raises some other questions. Suppose that, in

Case Four, the facts are the same as in *Case Three*, except that there is only one way in which you could prevent the nuclear explosion. You would have to start some fire that would cause the reactor to shut down. This fire would threaten to kill five people who would not be killed by the explosion because they are in an underground chamber. If you also start some flood, this flood would put out this fire after it has prevented the explosion, thereby saving the lives of these five people. This flood would also drown one other person who is in a deeper chamber. With these two acts, however, you would save the million people's lives in the way that would kill the fewest other people.

This case could take two forms. Suppose first that, to start the fire, you would have to press button A, which would also start the flood. This act would clearly be justified, though it would kill one person, because this act is the only way in which anyone could save the million people's lives.

Suppose next that, instead of pressing button A, which would start both the fire and the flood, you could press button B, which would start only the fire. You ought to press one of these buttons, since these are the only ways of saving the million people. The Time-Relative Harm Principle here implies that you ought to press button A, since that is the present possible act with which, in preventing the explosion, you would kill the fewest people.

It might be objected that I have here misapplied this principle. On this view, since both of your possible present acts would start the fire that would prevent the explosion, your choice between these acts ought to depend on the difference between them. This difference is that, if you start the flood as well as the fire, you would save five people's lives in a way that would also kill one other person. This act would be like your act, in *Hand Grenade*, if you stopped the runaway train that threatens to kill five people by throwing a bomb that would kill one other person. The Harm Principle implies that such acts are wrong. On this principle, you ought to start the fire, since that is the only way in which the million people's lives could be saved. But it would be wrong for you to start the flood, since you would thereby save only five people's lives in a way that would kill someone else. You ought to press button B, starting only the fire.

These claims are, I believe, mistaken. Suppose that, in

Case Five, lightning has struck some forest, starting a fire that now encircles, and threatens to kill, five people. You could save these people by starting a flood that would put out this fire, but this flood would kill one other person.

The Harm Principle does here imply that it would be wrong for you to start this flood, thereby saving only five people in a way that would kill one other person. If you did nothing, you would not save these five people's lives, but you would also kill no one. These claims do not apply, however, to *Case Four*. In this case it would be you, not a stroke of lightning, who would start the fire that would both prevent the explosion and threaten to kill five people. If you also start the flood, you would not be merely saving these five people's lives. This act would be your only way to avoid killing these people. Given this difference between *Cases Four* and *Five*, the Time-Relative Harm Principle rightly implies that, in *Case Four*, you ought to press button A, starting both the fire and the flood. That is how you could prevent the explosion, as you clearly ought to do, in the way that would kill the fewest people.

Suppose now that, in

Case Six, the facts are the same as in *Case Four*, except that there is no button A. You could start both the fire and the flood by simultaneously pressing buttons B and C.

The Time-Relative Harm Principle rightly implies that you ought to press both these buttons, thereby preventing the explosion in the way that kills the fewest people. But suppose next that, in

Case Seven, you cannot press these buttons at the same time. You could start the fire by pressing button B. You could then start the flood by pressing button C.

The Time-Relative Principle here implies that it would be wrong to press button C. If you pressed this button, thereby starting the flood, you would save only five people's lives, in a way that killed one other person. It is true that, by starting the flood, you would cause your earlier act of starting the fire not to kill these five people. On the Time-Relative Principle, however, you have a prima facie duty to act in the way that, of your possible *present* acts, would kill or harm the fewest people. It would now be wrong for you to start the flood because this act would kill one person, and, if you don't start the flood, your present act would kill no one.

These claims are clearly false. It would make no moral difference whether you press buttons B and C simultaneously or at different times. As the Temporally Neutral Harm Principle rightly implies, you ought first to start the fire and then later start the flood, since that is how your acts at these different times would together save the million people in the way that is best, because it would kill the fewest people.

It may seem that, rather than turning to this Temporally Neutral Principle, we could revise the Time-Relative Principle. Many of our acts take some time, and are complex acts that consist of a series of more particular acts. One such complex act might be your act of first starting the fire and then starting the flood. The Time-Relative Harm Principle might be widened so that it permits this complex act, which would prevent the explosion in the way that killed the fewest people. Suppose, however, that when you start the fire, you don't know that you could also start the flood, thereby causing your starting of the fire to kill four fewer people. It is only later that you realize that your starting the flood would have these effects. In this version of *Case Seven*, your starting the fire and starting the flood could not be claimed to be two parts of a single, complex act. These would be two quite different acts. To explain why you could justifiably first start the fire and then later start the flood, we should appeal to the Temporally Neutral Principle. These acts are justified because they are the acts at different times with which you could prevent the explosion in the way that kills the fewest people.

We can now compare, not time-relative and temporally neutral personal principles, but personal and omnipersonal principles. I shall continue to discuss cases that involve killing and saving lives. But it may help to note first how these distinctions apply to some other duties, such as our duties not to deceive or coerce people. These duties are often assumed to be personal, giving us the aims that we ourselves do not deceive or coerce people. But when we are considering cases in which no one acts wrongly, we could plausibly regard these duties as omnipersonal. Suppose that you and I both know some fact which ought to be kept secret from most other people, because this fact's becoming known would endanger many people's lives. I know that, unless I keep this fact secret by deceiving one person, it will become your duty to deceive several people. Would it be wrong for me to deceive one person, with the aim of preventing you from having a duty to deceive more people? Or suppose I know that, unless I coerce one person, it will become your duty to coerce several people. Would it be wrong for me to coerce this person, to prevent your having a duty to coerce more people?

These questions have been little discussed. People often claim that it would be wrong to deceive or coerce someone merely to minimize the number of people who are *wrongly* deceived or coerced by others. But such claims leave it open whether it would be wrong for us to deceive or coerce someone when we know, that if we don't do so, we would make it someone else's duty to deceive or coerce more people. If such acts would not be wrong, but would be morally required, our duties not to deceive or coerce people would not be personal in either of the senses that I have described. We would have these duties to everyone, and these duties would give everyone the common aims that we and others would justifiably deceive and coerce as few people as possible.

We can next compare two similar versions of a principle about saving people's lives. According to

the Personal Life-Saving Principle: We have a prima facie duty to do what would make it true that *we ourselves* save the most people.

According to

the Omnipersonal Principle: We have a prima facie duty to do what would make it true that *we and others* save the most people.

Suppose that, in

Case Eight, you could either save one person or act in a way that would enable me to save two or more other people.

As before, there are some versions of this case in which we could plausibly believe that you ought to save the single person. This might be true because you are a doctor and this person is one of your present patients, to whom you have special obligations. Or this person might be in front of you pleading for you to save her life, as she knows you could easily do. But if the people whose lives are in danger are all strangers to you and at some distance from you, it would be wrong for you to save one person rather than enabling me to save two or more other people. Suppose, for example, that in

Bicycle, there are three people whose lives are in danger some distance away. You have the knowledge that would be needed to save one of these people, and I have the knowledge that would be needed to save the other two people. There is only one way in which either you or I could reach any of these people soon enough to save their lives. You have a bicycle, which only one of us could use. If you used your bicycle, you would be able to reach and save the single person. If instead you let me use your bicycle, I would be able to reach and save the other two people.

We can plausibly believe that, in such cases, our duty to save people's lives would not be personal but omnipersonal. You could justifiably fail to save the single person's life because you would thereby enable me to save two other people's lives. We can make a stronger claim. In most cases of this kind, we would be making a serious moral mistake if we believed that we had a duty to save some stranger's life even when we would thereby fail to enable other people to save more people's lives. In such cases, we ought to have the common aim that we and others save as many lives as possible.

We can make a stronger claim. It would be wrong for you *not* to enable me to save these two other people. We would be making a serious moral mistake if we believed that we had a duty to save some stranger's life even when we know that, if we didn't do what would save this person's life, we would enable other people to save more people's lives. Our duty to save such people's lives, and to save them from lesser harms, are not in Ross's sense personal. We ought to have the shared or common aim that as many people as possible will be saved from death or lesser harms.

We can next compare two versions of the Harm Principle. According to

the Personal Harm Principle: We have a prima facie duty to do what would make it true that *we ourselves* kill or harm the fewest people.

According to

the Omnipersonal Harm Principle: We have a prima facie duty to do what would help to make it true that *we and others* kill or harm the fewest people.

The Personal Principle is most plausible when applied to certain cases in which some people act wrongly. Suppose that, in

Case Nine, some wrong-doers credibly threaten that, unless we kill some innocent person, they will kill two or more innocent people.

We can plausibly believe that we ought not to give into such threats. There are, we should admit, some exceptions. We ought to kill one innocent person if we would thereby persuade some wrong-doer not to explode some nuclear weapon that is hidden in some large city, thereby killing as many as a million people. But we can also plausibly believe that it would be wrong to kill one innocent person even if we would thereby persuade some wrong-doer not to carry out some threat to kill some much smaller number of people, such as five or ten.

Cases that involve such threats raise difficult questions. But these questions are irrelevant here, since we are asking what we ought to do in cases in which no one acts wrongly. We can add that, in cases like *Nine*, there is less disagreement between Act Consequentialism and Common Sense Morality. Act Consequentialists would believe that it would often be wrong to give into such threats by wrong-doers, even if we would thereby save a few people's lives, because giving into such threats would make wrong-doers more likely to make other such threats.

Suppose next that, in

Case Ten, you and I are the only people who could prevent the explosion that would kill a million people. You could prevent this explosion now, by doing something that would also kill one other person. If you do nothing, I would have to prevent this explosion with an act that would kill two other people.

On the Personal Harm Principle, you ought not to prevent the explosion, since this act of yours would kill one person, and if you did nothing you would kill no one. This principle is, I believe, mistaken. You could justifiably prevent this explosion in a way that would kill one person, since you know that, if you don't act in this way, you would make it my duty to prevent the explosion in a way that would kill two people. As the Omnipersonal Principle implies, in cases in which no one acts wrongly, we should have the common aim that we and others kill as few people as possible.

We can again revise this example, so that it raises some other questions. Suppose that, in

Case Eleven, you are in the control room of the damaged nuclear reactor. You have already pressed button B, thereby starting the fire that will

prevent the nuclear explosion and save the lives of a million people. This fire now threatens to kill five other people. You intend to press button C, which would start the flood that would put out this fire, thereby saving these five people, though in a way that would kill one other person. You find, however, that you cannot press button C, which is too high for you to reach. I am also in this control room, so you ask me to press this button on your behalf.

The Personal Harm Principle implies that you ought to start this flood, since this act would make it true that you have prevented the explosion in the way that kills the fewest people. But this principle also implies that it would be wrong for *me* to start this flood. Since I didn't start the fire, I have neither prevented the explosion, nor done what threatens to kill five people. If I start the flood, I would merely save the lives of five people in a way that would kill one other person. The Personal Harm Principle condemns such acts. If I refuse to start the flood, as you ask me to do, I shall have killed no one. That is why, if you can't reach button C, it would be wrong for me to press this button on your behalf.

Should we accept these claims? Could it be true that, though you ought to press button C, it would be wrong for me to press this button your behalf? The answer, I believe, is no. I am not assuming here that, whenever someone ought to do something, anyone else could justifiably do the same thing, or bring about the same outcome, on the first person's behalf. In cases that involve personal duties to other particular people, this may not be true. Suppose that you and I each have a child whose life is in danger. If I ought to save my child rather than yours, but I cannot act in this way, it would be wrong for you to act in this way on my behalf, by saving my child rather than yours. You ought instead to save your child. But no such claim applies to the cases that we are now discussing. If you can't start the flood, because button C is too high for you to reach, I could justifiably press this button on your behalf, thereby causing your earlier act of starting the fire to kill five fewer people.

Suppose next that the facts are slightly different. You and I could both reach both buttons. To start the fire and the flood, as we both know, we must press these buttons for two minutes. We might decide to do that in either of two ways. Suppose first that, in

Case Twelve, you press button B for two minutes, thereby starting the fire, and I press button C for two minutes, thereby starting the flood.

The Personal Harm Principle here implies that I have acted wrongly. By pressing button C for two minutes, I have merely saved five people's lives in a way that would kill one other person. This principle condemns such acts. Suppose next that, in

Case Thirteen, you and I each press each button for one minute, thereby starting both the fire and the flood.

The Personal Harm Principle would *not* here condemn my acts. If I had not pressed button B for one minute, your act of pressing button B would not have started the fire, so your act would not have saved the lives of the million people. By pressing each button for one

minute, I have therefore done what I ought to do, since I have helped to save these very many people's lives.

As before, this principle's claim about *Case Twelve* is false. If you and I press these buttons and thereby start both the fire and the flood, we would be acting rightly, since we would save the million people in the way that kills the fewest other people. It cannot make a moral difference whether we each press one button for two minutes, or both press both buttons for one minute.

Defenders of the Personal Harm Principle might now suggest that this principle should be revised. When different people cooperate in trying to achieve some aim, we can regard these people as if they were a single person, or a *collective agent*, and we should apply the Harm Principle to what these people *together* do. That is why, in *Cases Twelve* and *Thirteen*, it makes no moral difference whether we each press one button for two minutes, or both press both buttons for one minute. Whichever we do, we can claim that we together will have saved the million people in the best way that kills the fewest other people. Similar remarks apply to *Case Eleven*, in which you have pressed button B, but you cannot reach button C, and you ask me to press this button on your behalf. This revised principle would imply that, since I would be pressing button C on your behalf, we would again be a collective agent, who would together start both the fire and the flood. On this

Collective Harm Principle: We have a prima facie duty to do what would make it true that we and others would together kill or harm the fewest people.

This view implies that, even if you start the fire, and I merely start the flood that puts out this fire, I would be acting rightly. You and I would together save the million people in the way that is best, because it kills the fewest people.

This Collective Principle succeeds when applied to these cases. But this principle is, I believe, mistaken. Return first to cases in which we could save people's lives in ways that would kill no one. According to what we can call

the Collective Life-Saving Principle: We have a prima facie duty to do what would make it true that we and others would together save the most people.

Suppose that, in

Case Fourteen, many miners are trapped underground, with flood waters rising. There are two ways in which you and I could rescue some of these people. There is a large lift, which can be operated only by two people. We could together operate this lift, thereby saving 100 people. We could each instead separately operate one of two smaller lifts, with which we could each save 55 people.

The Collective Life-Saving Principle here implies that we ought to operate the large lift, since that is how we together would save the most people. If instead we operate the smaller lifts, we together would save no one. This principle is clearly mistaken. If we operate the large lift, we would together save only 100 people. If instead we operate the smaller lifts,

our separate acts would save the higher total of 110 people. In such cases, we should not have the common aim that we and others would *together* save the most people. We should have the common aim that we and others would in toto save the most people.

Suppose next that, in

Case Fifteen, after you press button B, thereby starting the fire, you become unconscious. I enter the control room, and I know that, if I now start the flood, I would cause your starting of the fire to prevent the explosion in the way that would kill the fewest people. For that reason, I press button C.

We cannot here appeal to the Collective Harm Principle, since you and I would not be acting together. We can suppose that you don't even know that I exist. But we can appeal to another, better principle. According to

the Omnipersonal Principle: We have a prima facie duty to do what would make it true both that we and others would in toto save the most people and that we and others would in toto kill the fewest people.

I ought to start the flood because I would thereby make it true that

(3) you and I save a total of a million people in a way that kills a total of only one person.

On this principle, as applied to *Case Fifteen*, it would be irrelevant that my act did not help you to save any of the million people, since you prevented the explosion by starting the fire all on your own. This principle appeals to the combined effects of our acts. By starting the flood I saved five people from being killed by your fire, and I also lowered the total number of people whom our acts would kill from five to one. If we appealed to either the Personal or Collective Principles, we would be led to ignore the fact that my starting of the flood would cause four fewer people to be killed. We should not ignore this fact, and we should therefore appeal to the Omnipersonal Principle.

Suppose next that, in

Case Sixteen, the facts are the same as in *Case Fifteen*, except that the fire was started, not by you, but by the nuclear reactor's automatic safety system.

It may seem that we cannot here appeal to the combined effects of different people's acts. But that may not be true. The people who installed this automatic safety system, we can suppose, intended this system to start a fire if that would be needed to prevent a nuclear explosion. The people who installed this system should then be regarded as being causally responsible for this fire. By starting the flood, I would cause it to be true that I and the installers of this safety system would in toto save the million people in the best way that would kill the fewest other people. The Omnipersonal Principle would then rightly permit this act.

This case may differ little from *Case Fifteen*, in which you start the fire, since you may be the person who installed the safety system, and you may have done that only yesterday. In another version of this case, however, the system was installed many years ago, and the installer has since died. This fact would not, I believe, make any moral difference. It would still be true that I could justifiably start the flood so that I and the installer of the safety system would in order to kill the fewest people.

These imagined cases would be unlikely to occur. But there are many actual cases that are relevantly similar. When we save people's lives, we often cause it to be true that we and others would in toto kill fewer people. That is true whenever we save people whose lives were threatened by some other people's past acts. I have been considering cases in which we could save such people's lives only by doing what would kill one other person. But there are many simpler cases in which, by saving certain people's lives, we would only be making it true that we and others would in toto kill fewer people. Suppose for example that, in

Case Seventeen, you have started the fire that saves a million people but also threatens to kill five others, and I save these people by starting a flood that puts out the fire, and that kills no one else.

By starting this flood, I would not only save these people's lives, but also cause your starting of the fire to kill five fewer people.

There are many actual cases in which, if we save some people's lives, we would also cause it to be true that we and others would in toto kill fewer people. That would be true, as I have said, even when our own acts would kill no one. In *Case Seventeen*, I would make it true that your act of saving the million kills no one rather than killing five other people.

We can take, as our examples, some of the ways in which rich people use the energies released by coal and oil in ways that overheat the Earth's atmosphere. These acts will later threaten many poor people's lives, by increasing the number and severity of floods, fires, droughts, famines, epidemics of infectious diseases, and other such causes of harm. It will often be true in the next few centuries that, to save many people from these threats to their lives, we or our descendants would have to act in ways whose foreseen side effects would be to kill some other smaller numbers of people. That may be true, for example, when we or others cause floods to put out fires, or isolate some people with infectious diseases. Many large scale policies save many people's lives in ways that unavoidably kill a few other people.

If these acts and policies merely saved people's lives in ways that killed several fewer other people, both versions of the Harm Principle would condemn these acts. In many of these cases, however, we and others would not merely be saving a larger number of people's lives. We or others would have earlier caused the threats to these people's lives. When we or others later do what would in toto *save* more people's lives, each of us will have contributed to the killing of fewer people. In such cases, we could appeal to the Omnipersonal Harm Principle. In explaining why these acts are justified, we need not reject the view that our duty not to kill is stronger than our duty to save people's lives. We could claim that these acts would cause it to be true that we and others have in toto killed fewer

people, thereby better fulfilling our duty not to kill people. Similar claims would apply to many lesser harms.

These claims apply, we can note, even when these threats to other people's lives were not in part caused by our own past acts. In the next few centuries there will be many young people who would themselves have done nothing to cause the threats to other people's lives that will be produced by climate change. By saving some of these other people's lives, however, these young people would cause the acts of some of their parents or grandparents to kill fewer people.

I shall now sum up some of these claims. I have been discussing whether, as many people assume, our duties not to kill or harm people are much stronger, in the conflict-of-duty sense, than our duties to save people from death or other harms. This claim is often assumed to conflict with the view that we ought to do what would cause the fewest people to be killed or harmed. This disagreement, I have claimed, is less deep than many people assume it to be. These views do not conflict in many cases in which, if we save several people's lives, we would cause it to be true that we or others would in toto kill fewer people. In my simple imagined case, if you start the fire and I start the flood, I would not merely be saving five people's lives in a way that would kill one other person. I would be making it true that you and I would in toto kill fewer people. It would be similarly though less obviously true that, when we or our descendants save some other people from being killed by the effects of climate change, we or our descendants would make it true that those who earlier caused these effects, by overheating the atmosphere, would have killed fewer people.

Here is another way to sum up these claims. If we earlier accepted some version of the Personal Harm Principle, and we responded to these arguments by coming to accept the Omnipersonal Principle, we would not be rejecting the widely held view that our duty not to harm people is much stronger than our duty to save people from harm. But this change in our beliefs would make this view less important. There would be many more cases in which we did not need to compare the strength of these duties, because these duties would not conflict. That would be true whenever, by saving some people's lives, we would also make it true that we and others would in toto kill fewer people. One example is the case in which, by starting the flood, I would not only save several people's lives, but would also cause your starting of the fire not to kill these people.

My imagined cases are in one way different from many actual cases. In my cases you would have a duty to start the fire which would prevent the explosion, thereby saving a million people. Most of us have no duty to act in the ways with which we are overheating the Earth's atmosphere. As we learn about the effects of these acts, we are coming to have duties *not* to act in these ways. But that does not affect the fact that, in saving some people from being killed by the effects of climate change, we would be making it true that we and others would in toto kill fewer people. These may be cases in which we and others have earlier acted wrongly. But this difference does not, I believe, undermine my main claims about the Personal and Omnipersonal Principles. Our descendants could justifiably make it true that our earlier wrong acts would kill fewer people.

There are some cases, however, in which it would make no difference whether we appeal to a Personal or to an Omnipersonal Principle. Suppose that, in

Case Eighteen, the facts are the same as before. You intend to start the fire that will prevent the explosion, thereby saving the million people. You intend later to start the flood, thereby saving the five people who would be killed by the fire, but killing one other person. Just before you act, however, some cat walks across the control panel and nudges button B, thereby starting the fire. In another version of this case, just before you act, a stroke of lightning starts the fire.

We cannot here claim that, if you press button C, thereby starting the flood, you would make it true that you saved the million people in the way that is best, because it killed the fewest people. Since you didn't start the fire, you have neither done what saved the million people, nor done what threatens to kill five other people. It was the cat, or the stroke of lightning, who or which did these things. We cannot claim that, if you now press button C, thereby starting the flood, you would cause your acts to kill four fewer people. You would be merely saving five people in a way that killed one other person. This act would be claimed to be wrong not only by the Personal Harm Principle, but also by the Omnipersonal Harm Principle.

We may find it hard to believe that it would now be *wrong* for you to do what you could have rightly done, one minute ago, if the cat or the stroke of lightning had not started the fire. By starting the flood, you would put out the fire, and thereby cause four fewer people to be killed. It may seem to make no moral difference whether it was you, or the cat, or the stroke of lightning that started the fire. We might be tempted to say that, in starting the flood, you wouldn't be merely saving five people in a way that would kill someone else. You would also be making it true that *you and the cat*, or *you and the stroke of lightning*, would in toto kill fewer people. But this claim would abandon the Omnipersonal Principle. If we believed that you could rightly start the fire, thereby causing four fewer people to be killed, we would be assuming that it made no moral difference whether these people would be intentionally killed by us, who are moral agents, or would instead be unintentionally killed by some non-moral event, such as the movement of a cat, or a stroke of lightning. On this view, which we can call

the Impersonal Principle: We have a prima facie duty to do what would make it true that the most people are saved, or that the fewest people are killed.

If we respond to *Case Eighteen* by appealing to this Impersonal Principle, we may seem to be rejecting my earlier claim that we ought to reject the Personal Harm Principle, and accept the Omnipersonal Principle. But this response would, I believe, be mistaken. There are many cases in which, in explaining what we ought to do, and why, we could justifiably appeal either to the effects of single acts, or to the effects of sets of acts. In some of these cases, we can more plausibly appeal to the effects of sets of acts. We should not assume that, if there are some cases in which we cannot appeal to such effects, we should never appeal to such effects.

One example is the imagined case that I called the *Harmless Torturers*. In the *Bad Old Days*, each of a thousand torturers turned some switch a thousand times on some machine, thereby inflicting great pain on one of a thousand victims. Things have now changed. Each

of these torturers now turns the switch once on each of the thousand machines. These acts together inflict the same great pain on each of the thousand victims. Because the effects of each act are so dispersed, however, none of the torturers now inflicts any perceptible increase of pain on any of the thousand victims.

I believe that, to explain the wrongness of these acts, we could appeal to the effects of each act. We could claim that each of the thousand torturers inflicts the same amount of pain on these victims, even though none of their acts makes any perceptible difference to the pain of any of the victims. We could more plausibly claim, however, that these torturers act wrongly because they together inflict the same amount of pain on all these victims. Similar remarks apply to cases in which there is over-determination. One example is that of firing squads, in which some group of soldiers together kill some victim, though each soldier could truly claim that, if he had disobeyed the order to fire, his failure to fire would have made no difference. Similar claims apply, and in the clearest and most important way, to many of the acts with which we and others are overheating the atmosphere. Though I believe that we could here appeal to the effects of each single act, most of us would find it more plausible to appeal to the ways in which the acts of us and others would later kill many people.

In *Case Eighteen*, however, in which the fire is started by some cat, or a stroke of lightning, we could not appeal to the combined effects of sets of acts.

[unfinished]