

3 Five Mistakes in Moral Mathematics

It is often claimed that, in cases that involve very many people, any single altruistic choice would make no difference. Some of those who make this claim believe that it undermines only the fourth moral solution, that provided by sufficient altruism. These people argue that, in such cases, because we *cannot* appeal to the *consequences* of our acts, we must appeal instead either to the Kantian Test, ‘What if everyone did that?’, or to the reluctance to free-ride.³⁹ But if my contribution involves a real cost to me, and would certainly make no difference—would give no benefit to others—I may not be moved by my reluctance to free-ride. This reluctance may apply only when I believe that I am profiting at the expense of others. If my contribution would make no difference, my failure to contribute will not be worse for others, so I will not be profiting at their expense. I may believe that the case is like those where some threshold has been clearly passed, so that any further altruistic act is a sheer waste of effort. This belief may also undermine the Kantian solution. If my contribution would make no difference, I can rationally will that everyone else does what I do. I can rationally will that no one contributes when he knows that his contribution would make no difference. Since others may think like me, it is of great importance whether, in such cases, we can explain why we should contribute by appealing to the consequences of our acts.

25. The Share-of-the-total View

We can. But to do so we must avoid several mistakes in moral mathematics. Consider

The First Rescue Mission: I know all of the following. A hundred miners are trapped in a shaft with flood-waters rising. These men can be brought to the surface in a lift raised by weights on long levers. If I and three other people go to stand on some platform, this will provide just enough weight to raise the lift, and will save the lives of these hundred men. If I do not join this rescue mission, I can go elsewhere and save, single-handedly, the lives of ten other people. There is a fifth potential

rescuer. If I go elsewhere, this person will join the other three, and these four will save the hundred miners.

When I could act in several ways, how should I decide which act would benefit people most? Suppose first that all five of us go to save the miners. On the *Share-of-the-Total View*, each produces his share of the total benefit. Since we five save a hundred lives, each saves twenty lives. Less literally, the good that each does is equivalent to the saving of this many lives. On this view, I ought to join the other four, and save the equivalent of twenty lives. I should not go elsewhere to save the other ten people, since I would then be saving fewer people. This is clearly the wrong answer. If I go with the other four, ten people needlessly die. Since the other four would, without my help, save the hundred miners, I should go and save these ten people.

The Share-of-the-Total View might be revised. It might be claimed that, when I join others who are doing good, the good that I do is not just my share of the total benefit produced. I should subtract from my share any reduction that my joining causes in the shares of the benefits produced by others. If I join this rescue mission, I shall be one of five people who together save a hundred lives. My share will be twenty lives. If I had not joined, the other four would have saved the hundred, and the share saved by each would have been twenty five lives, or five more than when I join. By joining I reduce the shares of the other four by a total of four times five, or twenty lives. On the revised view, my share of the benefit is therefore twenty minus twenty, or nothing. I should therefore go and save the other ten people. The revised view gives the right answer.

Consider next

The Second Rescue Mission. As before, the lives of a hundred people are in danger. These people can be saved if I and three other people join in a rescue mission. We four are the only people who could join this mission. If any of us fails to join, all of the hundred people will die. If I fail to join, I could go elsewhere and save, single-handedly, fifty other lives.

On the Revised Share-of-the-Total View, I ought to go elsewhere and save these fifty others. If instead I join this rescue mission, my share of the benefit produced is only the equivalent of saving twenty-five lives. I can therefore do more good if I go elsewhere and save the fifty. This is clearly false, since if I act in this way fifty more lives will be lost. I ought to join *this* rescue mission. We must make a further revision. I must add to my share of the benefit produced any increase that I cause to the shares produced by others. If I join, I enable each of three people to save, with me, a hundred lives. If I do not join, these three will save no lives. My share is twenty-five lives, and I increase by seventy-five the shares produced by the others. On this doubly revised view, my total share is a hundred lives. This is also the total share produced

by each of the others. Since each counts as producing the *whole* of this total benefit, this is not a version of the *Share-of-the-Total View*. It is a quite different view. This doubly revised view gives the right answer in this case. It is no objection to this view that it claims that each saves a hundred lives. This is what each does, not by himself, but with the help of the others.⁴⁰

This view can be put more simply. I should act in the way whose consequence is that most lives are saved. More generally,

(C6) An act benefits someone if its consequence is that someone is benefited more. An act harms someone if its consequence is that someone is harmed more. The act that benefits people most is the act whose consequence is that people are benefited most.

These claims imply, correctly, that I should not join the First Rescue Mission, but should join the Second.

Consequentialists should appeal to (C6). So should others, if they give any weight to what Ross called the *Principle of Beneficence*. On any plausible moral theory, we should sometimes try to do what would benefit people most.

(C6) may need to be further explained. Suppose that I can do either (1) or (2). In deciding which would benefit people more, I should compare *all* of the benefits and losses that people would later receive if I did (1), and *all* of the benefits and losses that people would later receive if I did (2). The act which benefits people more is the one that, in this comparison, would be followed by the greater *net* sum of benefits—the greater sum of benefits minus losses. It is irrelevant if, as is often true, the acts of many other people would also be parts of the cause of the receiving of these benefits and losses.

(C6) revises the ordinary use of the words ‘benefit’ and ‘harm’. When I claim to have benefited someone, I am usually taken to mean that some act of mine was the chief or immediate cause of some benefit received by this person. According to (C6), I benefit someone even when my act is a remote part of the cause of the receiving of this benefit. All that needs to be true is that, if I had acted otherwise, this person would not have received this benefit. Similar claims apply to ‘harm’.

There is a second way in which (C6) revises our use of ‘benefit’ and ‘harm’. On the ordinary use, I sometimes benefit someone even though what I am doing is not better for this person. This can be true when my act, though sufficient to produce some benefit, is not necessary. Suppose that I could easily save either J's life or K's arm. I know that, if I do not save J's life, someone else certainly will; but no one else can save K's arm. On our ordinary use, if I save J's life, I benefit him, and I give him a greater benefit than the benefit that I would give to K if I saved his arm. But, for moral purposes, this is not the way to measure benefits. In making my decision, I should ignore this benefit to J, as (C6) tells me to do. According to (C6), I do *not* benefit J when I save his life. It is not true that the consequence of my act is that J is benefited more. If I had acted differently, someone else would have

saved J's life. (C6) implies correctly that I ought to save K's arm. *This* is the act whose consequence is that people are benefited more. On the revised use of 'benefit', this is the act that benefits people more.

The First Mistake in moral mathematics is the Share-of-the-Total View. We should reject this view, and appeal instead to (C6).

26. Ignoring the Effects of Sets of Acts

It is natural to assume

(The Second Mistake) If some act is right or wrong *because of its effects*, the only relevant effects are the effects of this particular act.

This assumption is mistaken in at least two kinds of case.

In some cases, effects are *overdetermined*. Consider

Case One. X and Y simultaneously shoot and kill me. Either shot, by itself, would have killed.

Neither X nor Y acts in a way whose consequence is that an extra person dies. Given what the other does, it is true of each that, if he had not shot me, this would have made no difference. According to (C6), neither X nor Y harms me. Suppose that we make the Second Mistake: we assume that, if an act is wrong because of its effects, the only relevant effects are the effects of this particular act. Since neither X nor Y harms me, we are forced to the absurd conclusion that these two murderers do not act wrongly.

Some would take this case to show that we should reject (C6). There is a better alternative. We should add

(C7) Even if an act harms no one, this act may be wrong because it is one of a *set* of acts that *together* harm other people. Similarly, even if some act benefits no one, it can be what someone ought to do, because it is one of a set of acts that together benefit other people.

X and Y act wrongly because *they together* harm me. They together kill me. (C7) should be accepted even by Non-Consequentialists. On any plausible moral theory, it is a mistake in this kind of case to consider only the effects of single acts. On any plausible theory, even if each of us harms no one, we can be acting wrongly if we together harm other people.

In Case One, the overdetermining acts are simultaneous. What should we claim in cases where this is not true? Consider

Case Two. X tricks me into drinking poison, of a kind that causes a painful death within a few minutes. Before this poison has any effect, Y kills me painlessly.

Though Y kills me, Y's act is not worse for me. (C6) therefore implies that,

in killing me, Y does not harm me. (Y's act is in one way slightly worse for me, since it shortens my life by a few minutes. This is outweighed by the fact that Y saves me from a painful death.) (C6) also implies that X does not harm me. As in Case One, neither X nor Y harms me. (C7) implies correctly that X and Y act wrongly because they together harm me. They together harm me because, if *both* had acted differently, I would not have died.

Though (C7) gives the right answer here, this case may seem to provide an objection to (C6). It may seem absurd to claim that, in killing me, Y is not harming me. But consider

Case Three. As before, X tricks me into drinking poison of a kind that causes a painful death within a few minutes. Y knows that he can save *your* life if he acts in a way whose inevitable side-effect is my immediate and painless death. Because Y also knows that I am about to die painfully, Y acts in this way.

(C6) implies that Y ought to act in this way, since he will not harm me, and he will greatly benefit you. This is the right conclusion. Since Y's act is not worse for me, it is morally irrelevant that Y kills me. It is also morally irrelevant that X does *not* kill me. (C6) implies correctly that X acts wrongly. Though X does not kill me on the ordinary use of 'kill', he is here the real murderer. X harms me, and acts wrongly, because it is true that, if he had not poisoned me, Y would not have killed me. If X had acted differently, I would *not* have died. Y does not harm me because, if Y had acted differently, this would have made no difference to whether I died. Since Y does not harm me, and he greatly benefits you, Y is doing what he ought to do.

As these remarks show, Case Two provides no objection to (C6). In Case Three, (C6) correctly implies that Y ought to act as he does, because he does not harm me. In Case Two, Y's act affects me in just the same way. I was therefore right to claim that, in Case Two, Y does not harm me. Y acts wrongly in Case Two because he is intentionally a member of a group who together harm me.

It may be objected that, if *this* is why, in Case Two, Y acts wrongly, Y must be acting wrongly in Case Three. It may be thought that, here too, Y is intentionally a member of a group who together harm me.

This objection shows the need for another claim. In Case Three it is true that, if both X and Y had acted differently, I would not have been harmed. But this does not show that X and Y together harm me. It is also true that, if X, Y, *and Fred Astaire* had all acted differently, I would not have been harmed. But this does not make Fred Astaire a member of a group who together harm me. We should claim

(C8) When some group together harm or benefit other people, this group is the *smallest* group of whom it is true that, if they had all acted differently, the other people would not have been harmed, or benefited.

In Case Three, this ‘group’ consists of X. It is true of X that, if he had acted differently, Y would have done so too, and I would not have been harmed. No such claim is true of Y. In Case Two it is not true of either X or Y that, if he had acted differently, I would not have been harmed. I would not have been harmed only if *both* had acted differently. I would also not have been harmed if X, Y, and Fred Astaire had acted differently. But (C8) rightly implies that Fred Astaire is not a member of the group who together harm me. This group consists of X and Y.

Consider next

The Third Rescue Mission. As before, if four people stand on a platform, this will save the lives of a hundred miners. Five people stand on this platform.

Given what the others do, it is true of each of these five people that his act makes no difference. If he had not stood on this platform, the other four would have saved the hundred miners. Though none, by himself, makes any difference, these five together save the hundred miners. This case shows the need to add some further claim to (C8). In this case there is not *one* smallest group who together save the hundred lives. I shall return to such cases in Section 30.

There is a second kind of case in which we should consider the effects of sets of acts. These are *co-ordination problems*. One example is shown below.

		You	
		do (1)	do (2)
I	do (1)	Second-best	Bad
	do (2)	Bad	Best

Suppose that we apply Consequentialism only to single acts. We shall then claim that each has successfully followed C if he has done the act, of those that are possible for him, whose consequence is the best outcome. As we saw earlier, in co-ordination problems C will then be *indeterminate*. In this case we successfully follow C *both* if we both do (2) *and* if we both do (1).

Suppose that we both do (1). Given what you have done, I have done the act whose consequence is best. The outcome would have been worse if I had done (2). The same claims apply to you. If we both do (1) both successfully follow C, but we have not produced the best possible outcome.

Consequentialists should claim

(C9) Suppose that each of us has made the outcome as good as he can, given what the others did. Each has then acted rightly. But we together may have acted wrongly. This will be so if we together could have made the outcome better.

This is a claim about *objective* rightness, or what will in fact make the outcome best. If C includes this claim, it ceases to be indeterminate in this kind of case. (When we are deciding what to do, we should ask instead what is *subjectively* right, or what will be likely, given our beliefs, to make the outcome best. In most coordination problems of this kind, it is subjectively right for each of us to aim at the best outcome, since this is what others are likely to do.)⁴¹

27. Ignoring Small Chances

Return now to those Prisoner's Dilemmas that involve very many people. It is often claimed that, in these cases, we cannot appeal to the consequences of our acts. This involves at least two more mistakes.

One concerns those cases where each altruistic act has an extremely small chance of producing some extremely great benefit. It is sometimes claimed that, below some threshold, extremely small chances have no rational or moral significance.

This mistake is often made in discussions of elections. Though an election is not a pure Prisoner's Dilemma, it can illustrate this mistake. Many writers claim that, in a nation-wide election, one cannot justify voting merely by appealing to the consequences of one's act. But this is often false. Suppose that, if I vote, this will involve some costs, and will bring no benefits apart from the possible effect on who wins the election. On these assumptions, my voting cannot be justified in self-interested terms. But it can often be justified in Consequentialist terms. When I cannot predict the effects of my act, C tells me to do whatever would produce the greatest *expected* benefit. The expected benefit of my act is the possible benefit multiplied by the chance that my act will produce it. I may be able to justify my voting by appealing to this benefit.

Consider a Presidential Election in the United States. If I vote, there may be a very small chance that my vote will make a difference. On one estimate, if I am voting in one of the large and marginal states, which might go either

way, the chance that I shall make a difference would be about one in a hundred million. (The estimate is difficult. It should not be assumed that any pattern of votes is as likely as any other. But several writers agree that this chance is about one in a hundred million.⁴⁵)

Call the two candidates *Superior* and *Inferior*. And suppose that, if the next President is Superior, this will on average benefit Americans. There will be some Americans who will lose. It would have been better for these Americans if Inferior had won. But the losses to these Americans—the rich minority—will be outweighed by the benefits to all the other Americans. This is why Superior is the better candidate. If he is the one who is elected, this will produce a greater total net sum of benefits minus burdens. The average net benefit to Americans is this total sum divided by the number of Americans. For simplicity, I ignore effects on non-Americans. If my vote has a chance of one in a hundred million of affecting the result, the *expected* benefit of my voting is as shown below.

$$\frac{\text{The average net benefit to Americans from Superior's election} \times \text{the number of Americans}}{\text{One hundred million}} - \text{the costs to me and others of my voting}$$

Since there are two hundred million Americans, this sum is likely to be positive. This will be so if Superior's election would on average bring to Americans a net benefit more than half as great as the costs of my voting. I must be pretty cynical to doubt this. Similar remarks apply to many other public goods, and to altruists as well as Consequentialists. If an altruist does not ignore very tiny chances, he will often have a moral reason to make a contribution. The expected benefit that he would give to others would be greater than the costs of his contribution.

It may be objected that it is *irrational* to consider very tiny chances. When our acts cannot affect more than a few people, this may be so. But this is because the stakes are here comparatively low. Consider the risks of causing accidental death. It may be irrational to give any thought to a one-in-a-million chance of killing one person. But, if I was a nuclear engineer, would I be irrational to give any thought to the same chance of killing a million people? This is not what most of us believe. We believe, rightly, that such chances ought to be considered. Suppose that nuclear engineers did ignore all chances at or below the threshold of

one-in-a-million. It might then be the case that, for each of the many components in a nuclear reactor, there is a one-in-a-million chance that, in any day, this component would fail in a way that would cause a catastrophe. It would be clearly wrong for those who design reactors to ignore such tiny chances. If there are many reactors, each with many such components, it would not take many days before the one-in-a-million risk had been run a million times. There would fairly soon be a catastrophe.

When the stakes are very high, no chance, however small, should be ignored. The same is true when each chance will be taken very many times. In both these kinds of case, each tiny chance should be taken to be just what it is, and included in the calculation of the expected benefit. We can usually ignore a very small chance. But we should not do so when we may affect a very large number of people, or when the chance will be taken a very large number of times. These large numbers roughly cancel out the smallness of the chance.

A similar point applies if an act is likely or certain to give to others very small benefits. We should not ignore such benefits when they would go to a very large number of people. This large number roughly cancels out the smallness of the benefits. The total sum of benefits may thus be large.

These two points are not equally plausible. Very small benefits may be imperceptible. And it is plausible to claim that an ‘imperceptible benefit’ is *not* a benefit. But it is not plausible to claim that a very small chance is *not* a chance.

28. Ignoring Small or Imperceptible Effects

The Third Mistake in moral mathematics is to ignore very small chances when they would either affect very many people, or would be taken very many times. The Fourth and Fifth Mistakes are to ignore *very small* and *imperceptible* effects on very large numbers of people. These are similar mistakes, and can be criticised with the same arguments. But imperceptible effects raise one extra question.

I need not state both mistakes. The Fourth is the same as the Fifth except that ‘very small’ replaces ‘imperceptible’. Some people believe

(The Fifth Mistake) If some act has effects on other people that are imperceptible, this act cannot be morally wrong *because* it has these effects. An act cannot be wrong because of its effects on other people, if none of these people could ever notice any difference. Similarly, if some act would have imperceptible effects on other people, these effects cannot make this act what someone ought to do.

One kind of imperceptible effect is not controversial. I may cause you serious harm *in a way* that is imperceptible. The dose of radiation that I give you may

be the unknown cause of the cancer that kills you many years later. Though the cause may be unknown, the effect is here perceptible. But, in the cases I shall be considering, the *effects* are imperceptible.

Consider first a variant of a case described by Glover.⁴⁴

The Drops of Water. A large number of wounded men lie out in the desert, suffering from intense thirst. We are an equally large number of altruists, each of whom has a pint of water. We could pour these pints into a water-cart. This would be driven into the desert, and our water would be shared equally between all these many wounded men. By adding his pint, each of us would enable each wounded man to drink slightly more water—perhaps only an extra drop. Even to a very thirsty man, each of these extra drops would be a very small benefit. The effect on each man might even be imperceptible.

Assume that the benefit given to each man would be merely the relieving of his intensely painful thirst. There would be no effect on these men's health. Since the benefits would be merely the relief of suffering, these are the kind of benefit of which it can most plausibly be claimed that, to be benefits at all, they must be perceptible.

Suppose first that, because the numbers are not very large, the benefit that each of us would give to each man would, though very small, be perceptible. If we make the Fourth Mistake, we believe that such tiny benefits have no moral significance. We believe that, if some act would give to others such tiny benefits, this cannot make this act what someone ought to do. We are forced to conclude that none of us ought to add his pint. This is clearly the wrong conclusion.

Assume, next, that there are a thousand wounded men, and a thousand altruists. If we pour our pints into the water-cart, each of us will cause each wounded man to drink an extra thousandth of a pint. These men might notice the difference between drinking *no* water and one thousandth of a pint. Let us therefore ask, 'If these men will drink at least one tenth of a pint, could they notice the effect of drinking any extra thousandth of a pint?' I shall assume that the answer is No. (If the answer is Yes, we merely need to suppose that there are more altruists and wounded men. There must be some fraction of a pint whose effect would be too small to be perceptible.)

Suppose that a hundred altruists have already poured their water into the cart. Each of the wounded men will drink at least one tenth of a pint. We are the other nine hundred altruists, each of whom could add his pint. Suppose next that we make the Fifth Mistake. We believe that, if some act would have imperceptible effects on other people, these effects cannot make this act what someone ought to do. If we believe this, we cannot explain why each of us ought to add his pint.

It may be said: 'We can avoid this problem if we redescribe the effect of adding each pint. We need not claim that this gives to each of the men one

thousandth of a pint. We could claim that it gives to one man one pint.’

This claim is false. The water will be shared equally between all these men. When I add my pint, is the effect that an extra man receives a full pint? If I had not added my pint, is there some man who would have received nothing rather than a full pint? Neither of these is true. There is only one correct description of the effect of my act. It gives to each of the thousand men an extra thousandth of a pint.

It may next be said that we should appeal to the Share-of-the-Total-View. On this view, the share that each contributes is equivalent to the benefit that one man receives from one pint. But we cannot appeal to this view, since we saw in Section 25 that it can imply absurd conclusions.

What we can appeal to is a claim about what we *together* do. We can claim

(C10) When (1) the best outcome would be the one in which people are benefited most, and (2) each of the members of some group could act in a certain way, and (3) they would benefit these other people if *enough* of them acted in this way, and (4) they would benefit these people *most* if they *all* acted in this way, and (5) each of them both knows these facts and believes that enough of them will act in this way, then (6) each of them ought to act in this way.

Each of us could give to each of the thousand wounded men an extra thousandth of a pint of water. If enough of us act in this way this will benefit each of these men. And we will benefit these men most if we all act in this way. We know these facts, and we know that enough of us—one hundred—have already acted in this way. (C10) implies correctly that each of us ought to act in this way.

Remember now the Fifth Mistake. On this view, an act cannot be right or wrong, *because* of its effects on other people, if these effects are imperceptible. The case just described refutes this view. It is clear that, in this case, each of us should pour his pint into the water-cart. Each of us should cause each wounded man to drink an extra thousandth of a pint. Each of us ought to affect each wounded man in this way, even though these effects are imperceptible. We may believe that, because these effects are imperceptible, each of us is benefiting no one. But, even if *each* benefits no one, *we together* greatly benefit these wounded men. The effects of *all* our acts are perceptible. We greatly relieve the intense thirst of these men.

Consequentialists may appeal to several principles. They may thus believe that, in some cases, the best outcome is not the one in which people are benefited most. To cover such cases, they can claim

(C11) When (1) the members of some group would make the outcome better if *enough* of them acted in some way, and (2) they would make the outcome *best* if *all* of them acted in this way, and (3)

each of them both knows these facts and believes that enough of them will act in this way, then (4) each of them ought to act in this way.

Non-Consequentialists believe that, in some cases, we should try to produce the best outcome. In some of these cases, they can appeal to (C11). As before, in some cases, (C11) does not by itself give the right answer. We would have to add a more complicated claim. I shall ignore these complications.⁴⁶

As I showed in Section 26, there are two kinds of case where we need to appeal to the effects, not just of single acts, but of sets of acts. We need to make this appeal when (1) the effects of our acts are overdetermined, or (2) we face coordination problems. We are now considering cases where (3) each person's act will have imperceptible effects on other people. This may be a third kind of case in which we need to appeal to the effects of sets of acts. Whether we *need* to make this appeal depends in part on the answer to another question.

29. Can There Be Imperceptible Harms and Benefits?

It may be objected: 'You claim that each of the thousand altruists should pour in his pint, since this is how the wounded men would be benefited most. This claim is false. Suppose that one of the altruists does not pour in his pint. Are the wounded men benefited less? They are not. They drink slightly less water. But this effect is imperceptible. Since the effect is imperceptible, the benefit to these men cannot be less.'

This objection assumes that there cannot be imperceptible benefits. If we make this assumption, we face part of a wider problem, variously called the *Sorites Problem*, *Wang's Paradox*, or the *Paradox of the Heap*.

In our case, the benefit is the relieving of intensely painful thirst. If each man receives a pint of water, his thirst will become less painful. His pain will be less bad. Our problem is the following. We assume

(A) Someone's pain cannot become *imperceptibly* better or worse. Someone's pain cannot become either less bad, or worse, if this person could not possibly notice any difference.

And it is plausible to assume

(B) *At least as bad as* and *not worse than* are, when applied to pains, *transitive* relations. Thus, if someone's pain in Outcome (2) is at least as bad as it was in Outcome (1), and his pain in Outcome (3) is at least as bad as it was in Outcome (2), his pain in Outcome (3) must be at least as bad as it was in Outcome (1).

A hundred altruists have already poured in their pints. Each of the wounded men will drink at least a tenth of a pint. They would not notice the effect of any extra thousandth of a pint. In different possible outcomes different numbers of altruists will later pour their pints into the cart. Let us refer to these outcomes by citing the number who contribute. Thus, if no one else contributes, this will produce Outcome 100.

Suppose that one more altruist contributes. Each wounded man will drink more water, but the amount will be so small that he cannot notice this. According to (A), each man's thirst cannot become less painful. Each man's pain in Outcome 101 must be at least as bad as it would have been in Outcome 100. Suppose next that a second altruist adds his pint. As before, none of the men can notice this difference. According to (A) each man's thirst cannot become less painful. Each man's pain in Outcome 102 must be at least as bad it would have been in Outcome 101. According to (B) each man's pain in Outcome 102 must be at least as bad as it would have been in Outcome 100. The same claims apply if a third altruist contributes. In Outcome 103 each man's pain must be at least as bad as it would have been in Outcome 100. These claims apply to every extra altruist who contributes. Suppose that all of us contribute. The result is Outcome 1000, in which each man drinks a whole pint. (A) and (B) together imply that each man's pain must be at least bad as it would have been in Outcome 100. Drinking a whole pint, rather than only one tenth, cannot do anything to relieve the pain of each man's thirst. Since this conclusion is absurd, we must reject either (A) or (B).

Which should go? I reject (A). I believe that someone's pain can become less painful, or less bad, by an amount too small to be noticed. Someone's pain is worse, in the sense that has moral relevance, if this person minds the pain more, or has a stronger desire that the pain cease. I believe that someone can mind his pain slightly less, or have a slightly weaker desire that his pain cease, even though he cannot notice any difference. More generally, there can be imperceptible harms, and imperceptible benefits. In many other kinds of case, people have been shown to make very small mistakes when they report the nature of their experiences. Why should we assume that they cannot make such mistakes about the strength of their desire that some pain cease?

Suppose that you reject these claims, and continue to accept (A). You must then reject (B). To avoid the absurd conclusion reached above, you must admit that, when applied to pains, *at least as bad as* and *not worse than* are not transitive relations. And rejecting (B) has implications like those of rejecting (A). You must now admit that your acts may be wrong, because of their effects on someone's else's pain, even though *none* of your acts makes this person's pain worse. They may *together* have this effect. Each act may be wrong, though its effects are imperceptible, because it is one of a set of acts that together make this person's pain very much worse.

Consider

The Bad Old Days. A thousand torturers have a thousand victims. At the start of each day, each of the victims is already feeling mild pain. Each of the torturers turns a switch a thousand times on some instrument. Each turning of a switch affects some victim's pain in a way that is imperceptible. But, after each torturer has turned his switch a thousand times, he has inflicted severe pain on his victim.

Suppose that you make the Fifth Mistake. You believe that an act cannot be wrong, because of its effects on other people, if these effects are imperceptible. You must then conclude that, in this case, no turning of a switch is wrong. None of these torturers ever act wrongly. This conclusion is absurd.

Why are the torturers acting wrongly? One explanation appeals to the total effect of what each torturer does. Each turns a switch a thousand times. These acts, taken together, inflict severe pain on his victim.

Consider next

The Harmless Torturers. In the Bad Old Days, each torturer inflicted severe pain on one victim. Things have now changed. Each of the thousand torturers presses a button, thereby turning the switch once on each of the thousand instruments. The victims suffer the same severe pain. But none of the torturers makes any victim's pain perceptibly worse.

Can we appeal here to the total effect of what each torturer does? This depends in part on whether we reject (A), believing that someone's pain can become imperceptibly worse. If we believe this, we can claim: 'By pressing the button, each torturer causes each victim to suffer slightly more. The effect on each is slight. But, since each torturer adds to the suffering of a thousand victims, he imposes on them a great total sum of suffering. Since the victims suffer just as much as they did in the Bad Old Days, these torturers are acting just as wrongly as they used to do. In the Bad Old Days, each torturer imposed on one victim a great sum of suffering. Each of the Harmless Torturers imposes on these thousand victims an equally great total sum of suffering.'

Suppose instead that we accept (A), believing that pains cannot become imperceptibly worse. We must then admit that each of the Harmless Torturers causes no one to suffer more. We cannot appeal to the total effect of what each torturer does. On our view, none of the torturers harms anyone.

Even if none of them harms anyone, the torturers are clearly acting wrongly. If we cannot appeal to the effects of what each torturer does, we must appeal to what the torturers together do. Even if none of them causes any pain, they together impose great suffering on a thousand victims. We can

claim

(C12) When (1) the outcome would be worse if people suffered more, and (2) each of the members of some group could act in a certain way, and (3) they would cause other people to suffer if *enough* of them acted in this way, and (4) they would cause these people to suffer *most* if they *all* acted in this way, and (5) each of them both knows these facts and believes that enough of them will act in this way, then (6) each would be acting wrongly if he acted in this way.

Someone may again object: 'In the case of the Harmless Torturers, (4) is not true. These torturers do not cause their victims to suffer *most* if they *all* turn each switch once. Suppose that one of them turned no switches. None of the victims would notice any difference. Since a pain cannot become imperceptibly less bad, the victims would *not* suffer less if one of the torturers did not act.'

As I remarked, this objection raises the well-known Sorites Problem. If we accept (A), our answer to this objection must involve a solution to this problem. Since this problem is both hard to solve, and raises questions which have nothing to do with ethics, I shall not discuss it here.⁴⁸

If we accept (A), our objection to the Harmless Torturers must both be complicated and solve the Sorites Problem. If we reject (A), our objection could be simple. We could claim that each of the torturers inflicts on the victims a great total sum of suffering.

Of these two explanations, which is better? Even if we reject (A), we may be wrong to give the simpler explanation. Whether this is so depends on the answer to another question. Consider

The Single Torturer. One morning, only one of the torturers turns up for work. It happens to be true that, through natural causes, each of the victims is already suffering fairly severe pain. This pain is about as bad as it would be after the switches had been turned five hundred times. Knowing this fact, the Single Torturer presses the button that turns the switch once on all of the machines. The effect is the same as in the days when all the torturers act. More precisely, the effect is just like that when each switch is turned for the five hundred and first time. The Single Torturer knows that this is the effect. He knows that he is not making any victim's pain perceptibly worse. And he knows that he is not a member of a group who together do this.

Is the Single Torturer acting wrongly? Suppose we believe that he is not. We cannot then appeal to the simpler objection in the case where all the torturers act. We cannot claim that each is acting wrongly because he is imposing on others a great total sum of suffering. If this is why each is acting wrongly, the Single Torturer must be acting wrongly. He acts in the

same way, and with the same effects. If we believe that the Single Torturer is not acting wrongly, we must give the other objection in the case where all the torturers act. We must claim that each is acting wrongly because he is a member of a group who together inflict great suffering on their victims.

I believe that the Single Torturer *is* acting wrongly. How can it make a moral difference whether he produces bad effects jointly with other agents, or with Nature?⁴⁹ I therefore prefer, in both cases, to appeal to the effects of single acts.

Some people disagree. Even if we believe that there can be imperceptible harms and benefits, it may thus be better to appeal to what groups together do. This appeal is less controversial.

(If the Single Torturer is *not* acting wrongly, it may be unfair to claim that some of us make *five* mistakes in moral mathematics. On this view, the Fifth Mistake is merely a special case of the Second Mistake. But this would seldom make a practical difference.)

In this section I have asked whether there can be imperceptible harms and benefits. I am inclined to answer Yes. If we answer No, we must abandon the claim that, when applied to harms and benefits, *at least as bad as* and *not worse than* are transitive relations. I have also shown that it makes little difference which answer we accept. On either answer, we must abandon what I call the Fifth Mistake. We must abandon the view that an act cannot be either right or wrong, *because* of its effects on other people, if these effects are imperceptible.

30. Overdetermination

Return now to the pints of water and the wounded men. Let us add some features to this case. Suppose that, before the water-cart is driven to these men, you arrive, with another pint. The wounded men need more than a single pint. After drinking this pint their intensely painful thirst would not be fully relieved. But the water-cart can hold only one thousand pints. It is now full. If you add your pint, this will merely cause one pint to overflow down some drain.

You have no moral reason to add your pint, since this would merely cause a pint to be wasted. According to (C10), you ought to add your pint if this would make you a member of a group who together benefit other people. We may think that, if you add your pint, you are *not* a member of the group who together benefit the wounded men. Some of your pint may be drunk by these wounded men. And you are acting in the same way as the other altruists did. But we might claim, 'Unlike the other altruists, you do not give to each of the wounded men an extra thousandth of a pint of water. Your act has no effect on the amount of water that these men receive.'

Things are not so simple. If you add your pint, this will be a case

involving overdetermination. It is true of you that, if you had not contributed, this would have made no difference to the amount of water that the men drink. But, since you have contributed, the same is true of each of the other altruists. It is true that, if any one of these altruists had not contributed, this would have made no difference to the amount of water that the men drink. The water-cart would not have been full when you arrive, and your pint would have made it full. What is true of you is true of each of the other altruists. It is therefore true that you *are* a member of the group who together benefit the wounded men.

We must again appeal to what the agents know, or have reason to believe. Suppose that the other altruists had no reason to believe that you would arrive, with your extra pint. Each ought to have poured in his pint. This is because each had good reason to believe that he would be a member of a group of whom it is true both (1) that they together benefit the wounded men, and (2) that they benefit these men *most* if they *all* pour in their pints. When you arrive, you know that the water-cart is full. You have no reason to contribute, since you know that you would *not* be a member of such a group. If you contribute, you will instead be a member of a group *which is too large*. We should claim

(C13) Suppose that there is some group who, by acting in a certain way, will together benefit other people. If someone believes that this group either is, or would be if he joined, *too large*, he has no moral reason to join this group. A group is *too large* if it is true that, if one or more of its members had not acted, this would not have reduced the benefit that this group gives to other people.

If you add your pint, this will make this group of altruists too large. If you do *not* add your pint, this group will *not* be too large. This is a special borderline case. (C13) also covers the more common cases where some group is already too large.

31. Rational Altruism

The Fifth Mistake in moral mathematics is the belief that imperceptible effects cannot be morally significant. This is a very serious mistake. When all the Harmless Torturers act, each is acting *very* wrongly. This is true even though each makes no one perceptibly worse off. The same could be true of us. We should cease to think that an act cannot be wrong, *because* of its effects on other people, if this act makes no one perceptibly worse off. Each of our acts may be *very* wrong, because of its effects on other people, even if none of these people could ever notice any of these effects. Our acts may *together* make these people very much worse off.

The Fourth Mistake is equally serious. If we believe that trivial effects can

be morally ignored, we may often make people very much worse off. Remember the Fisherman's Dilemma. Where there is overfishing, or declining stocks, it can be better for each if he tries to catch more, worse for each if all do. Consider

How the Fishermen Cause a Disaster. There are many fishermen, who earn their living by fishing separately on some large lake. If each fisherman does not restrict his catch, he will catch within the next few seasons more fish. But he will thereby lower the total catch by a much larger number. Since there are many fishermen, if each does not restrict his catch, he will only trivially affect the number caught by each of the others. The fishermen believe that such trivial effects can be morally ignored. Because they believe this, even though they never do what they believe to be wrong, they do not restrict their catches. Each thereby increases his own catch, but causes a much greater lowering in the total catch. Because they all act in this way, the result is a disaster. After a few seasons, all catch very many fewer fish. They cannot feed themselves or their children.

If these fisherman knew the facts, had sufficient altruism, and avoided the Fourth Mistake, they would escape this disaster. Each knows that, if he does not restrict his catch, this will be somewhat better for himself, whatever others do. And each knows that, if he acts in this way, the effects on each of the others will be trivial. But the fishermen should not believe that these trivial effects can be morally ignored. They should believe that acting in this way is wrong.

As before, there are two ways in which we could explain why these acts are wrong. We could appeal to the total effect of each person's act. Each fisherman knows that, if he does not restrict his catch, he will catch more fish, but he will reduce the total catch by a much larger number. For the sake of a small gain to himself, he imposes on others a much greater total loss. We could claim that such acts are wrong. This claim does not assume that there can be imperceptible harms and benefits. It is therefore less controversial than the corresponding claim about what each of the Harmless Torturers does.

Our alternative is to appeal to what these fishermen together do. Each fisherman knows that, if he and all the others do not restrict their catches, they will together impose upon themselves a great total loss. Rational altruists would believe these acts to be wrong. They would avoid this disaster.

It may be said: 'So would rational egoists. Each knows that, if he does not restrict his catch, he is a member of a group who impose upon themselves a great loss. It is irrational to act in this way, even in self-interested terms.' As I shall argue in the next chapter, this claim is not justified. Each knows that, if he does not restrict his catch, this will be *better* for himself. This is so whatever others do. When someone does what he knows will be better for

himself, it cannot be claimed that his act is irrational in self-interested terms.

Remember next

The Commuter's Dilemma. Suppose that we live in the suburbs of a large city. We can get to and return from work either by car or by bus. Since there are no bus-lanes, extra traffic slows buses just as much as it slows cars. We could therefore know the following to be true. When most of us are going by car, if any one of us goes by car rather than by bus, he will thereby save himself some time, but he will impose on others a much greater total loss of time. This effect would be dispersed. Each might cause a hundred others to be delayed for twenty seconds, or cause a thousand others to be delayed for two seconds. Most of us would regard such effects as so trivial that they can be morally ignored. We would then believe that, in this Commuter's Dilemma, even a rational altruist can justifiably choose to go by car rather than by bus. But if most of us make this choice we shall all be delayed for a long time every day.

Rational altruists would avoid this result. As before, they could appeal either to the effects of what each person does, or to the effects of what all together do. Each saves himself some time, at the cost of imposing on others a much greater total loss of time. We could claim that it is wrong to act in this way, even though the effects on each of the others would be trivial. We could instead claim that this act is wrong, because those who act in this way together impose on everyone a great loss of time. If we accept either of these claims, and have sufficient altruism, we would solve the Commuter's Dilemma, saving ourselves much time every day.

Similar reasoning applies to countless other cases. For one more example, consider the devices that purify the gases that our cars emit. We would think it wrong to save ourselves the cost of repairing this device, if in consequence we imposed great air-pollution on some other single person. But many of us would not think this wrong if it merely trivially or imperceptibly increased the air-pollution suffered by each of very many people. This would be the actual effect in many large cities. It might be much better for all of us if none of us caused such pollution. But, to believe that we are acting wrongly, many of us need to change our view. We must cease to believe that an act cannot be wrong, because of its effects on other people, if these effects are either trivial or imperceptible.

As conditions change, we may need to make some changes in the way we think about morality. I have been arguing for one such change. Common-Sense Morality works best in small communities. When there are few of us, if we give to or impose on others great total benefits or harms, we must be affecting other people in significant ways, that would be grounds either for gratitude, or resentment. In small communities, it is a plausible claim that we cannot have harmed others if there is no one with an obvious complaint, or

ground for resenting what we have done.

Until this century, most of mankind lived in small communities. What each did could affect only a few others. But conditions have now changed. Each of us can now, in countless ways, affect countless other people. We can have real though small effects on thousands or millions of people. When these effects are widely dispersed, they may be either trivial, or imperceptible. It now makes a great difference whether we continue to believe that we cannot have greatly harmed or benefited others unless there are people with obvious grounds for resentment or gratitude. While we continue to believe this, even if we care about effects on others, we may fail to solve many serious Prisoner's Dilemmas. For the sake of small benefits to ourselves, or our families, each of us may deny others much greater total benefits, or impose on others much greater total harms. We may think this permissible because the effects on each of the others will be either trivial or imperceptible. If this is what we think, what we do will often be much worse for all of us.

If we cared sufficiently about effects on others, and changed our moral view, we would solve such problems. It is not enough to ask, 'Will my act harm other people?' Even if the answer is No, my act may still be wrong, because of its effects. The effects that it will have when it is considered on its own may not be its only relevant effects. I should ask, 'Will my act be one of a set of acts that will *together* harm other people?' The answer may be Yes. And the harm to others may be great. If this is so, I may be acting *very* wrongly, like the Harmless Torturers. We must accept this view if our concern for others is to yield solutions to most of the many Prisoner's Dilemmas that we face: most of the many cases where, if each of us rather than none of us does what will be better for himself—or for his family, or those he loves—this will be worse, and often *much* worse, for everyone.