

please cite published version (which contains printing errors):
In Michael J. Almeida (ed.), *Imperceptible Harms and Benefits*
(Dordrecht: Kluwer, 2000), pp. 9-42

POOLED BENEFICENCE

Garrett Cullity

I: The Problem of Pooled Beneficence

When many of us find that we share an aim, we tend to cooperate in pursuing it. This makes sense; for by cooperating, we will often collectively be able to achieve that aim more effectively than if we had each pursued it independently. When we do this, we turn our common aim into a collective aim — the aim of a collective action that we perform together.

This often happens with people who share a beneficent aim. When there is a large-scale need, to which many of us think we ought to respond, we tend to do so by pooling individual contributions and then organizing the distribution of resources from the pool, to the best overall effect. In this case, the advantages of cooperation are especially clear. If more people can be helped to a greater extent than by uncoordinated acts of beneficence, then this seems to provide us with a very clear case for cooperation. Indeed, it seems that there are often contexts in which the existence of collectively organized, large-scale beneficent activity creates opportunities for beneficence which would not otherwise exist. Where people who need help are distant from us, or a threshold of resources needs to be surpassed before the help can be effective, there may be little any of us could usefully do alone, whereas by combining our efforts we may be able to achieve a lot.

By this means, we may be able to achieve a lot collectively; but what does each of us achieve individually? In various ways, it can seem that the answer to this is: Nothing significant at all. For whether or not I contribute to the pool will make no perceptible difference to its beneficiaries.

Philosophical discussion of this has often focused on examples of pools for beneficence which have two features. First, they have a *limited capacity*, and secondly, each contribution produces an effect which is dispersed among all the beneficiaries of the pool — referring to this second feature, I shall say that they are *dispersing*. Thus, in Derek Parfit's well-known example of a pool for beneficence, we imagine that each of a thousand people could add a pint of water to a thousand-pint tank the contents of which will be distributed equally to a thousand severely thirsty men.¹ This is a limited-capacity pool: after a thousand pints it is full. And it is dispersing, since the effect of any individual's adding a pint will be that each of the men receives one thousandth of a pint more than he otherwise would have. In virtue of this second feature, it seems that no individual act of contributing to the pool will leave anyone perceptibly better off.

Notice, however, that actual pools for beneficence do not normally have either of these features. There is no point at which the fund of an international aid agency, for example, is full.² And although it is tempting to think of them as dispersing, this is often false as well. Thus consider an aid agency fund for famine relief. Suppose I donate enough money to feed one person for the duration of a crisis. My donation will certainly not be used to do this. However, it is tempting to think that it will be used to increase the agency's overall stock of food by this amount, and the increase will be dispersed among a large number of hungry people. On reflection, though, this looks wrong. Relief agencies do not buy food in such small amounts. Instead, they make a large-scale calculation of the size of the overall need, the amount of money they are likely to be able to raise from various sources to pay for it, and the extent to which it

makes sense to draw on contingency funds in the light of other likely calls on them. Guided by these large-scale considerations, they then arrange the shipment of a large quantity of food. If so, it seems wrong to think that my donation to a pool of this kind will produce an effect which is dispersed among the pool's beneficiaries; rather, it will have no effect at all on them. For it is plausible to hold that whether or not I make a donation will make no difference to the agency's deliberations: they are conducted on a scale that is simply not sensitive to the difference I would make by donating enough to help one person.

Now this may seem to invite the following reply. Perhaps it is true that my donation will *probably* have no effect on whether another food consignment is sent. But what is important is not the actual but the expected benefit associated with my action. And if there is a small chance that my action will produce a large benefit, then the expected benefit may still be significant. After all, if an agency's appeal fund continues to grow as individual contributions are made to it, then eventually it will send another consignment. And for each consignment sent by an efficiently-run agency, reliable expectations can be formed of the number of lives that will be saved by it. So the question is whether my donation is the one that triggers the threshold which causes the agency to send another consignment. If my donation is, say, one thousandth of the amount needed to make a consignment that can be expected to save a thousand lives, then on average my donating that amount would trigger that threshold once in every thousand times I made it. This gives me a one in a thousand chance that my donation will be the one that makes the difference between saving and not saving a thousand lives. So the expected benefit is the saving of one life.

However, what can plausibly be denied is that there is any such threshold which is triggered by a single donation. Aid agency decisions about whether to send another food consignment are certainly sensitive to the size of their appeal fund; and donors to

that fund will collectively affect the decision over whether to send more food. But those decisions will not be so sensitive to the exact size of the fund that any single donation will have such an effect; and if not, there is not even a small chance that my donation will be the one to have that effect. Given the agency's assessment of the amount of food it is desirable to buy, any small shortfall that results from my not making a donation will be covered — indeed, much larger shortfalls will be covered — by extra fundraising efforts, or absorbed in administrative savings.

If this is right, it gives us two kinds of pool for beneficence, and two different ways in which it whether or not I contribute will make no perceptible difference to beneficiaries of the pool. In cases of the first, dispersing, kind, my contributing to the pool does make a difference to its beneficiaries, but not a perceptible one. In the second, non-dispersing kind of case, it seems that whether or not I contribute will make no difference to them at all. The collective action that gets taken will not be affected at all by whether or not I join in. To these we can add a third kind of case, where it is reasonable for me to believe that if I do not contribute to a limited-capacity pool someone else will.³ For example, a rescue party of a certain size might be needed, and when I am asked to join I may know that if I refuse someone else will take my place. Here (unlike the first two cases), it seems that if I join in I do perceptibly benefit the people who are helped. The fact that someone else would have produced the effect if I had not does not show that I do not produce that effect. But it still seems true that it makes no perceptible difference to those saved by the rescue party whether or not I have decided to join. If I had not benefited them, someone else would have.

If, in these different ways, it will not make a perceptible difference to the beneficiaries of such pools whether or not I contribute, this can encourage the conclusion that it cannot be wrong for me not to contribute. Of course, if I affect you in a way that is not now perceptible, but which will lead later on to your being

perceptibly worse off — I give you a slow-acting poison, say — then there is not much temptation to say that what I do cannot be wrong. But if you will *never* be perceptibly better off as a result of what I am being asked to do, then it can seem attractive to hold that refusing to do it cannot be wrong. After all, whether or not I do it ought to be a matter of indifference to the intended beneficiaries. But, since it involves a perceptible cost, it will not be a matter of indifference to me.

Someone who endorses this argument I call, rather inelegantly, an "imperceptibilist". Imperceptibilists hold that if my contribution to a collective activity of beneficence will never make a perceptible difference to anyone, it cannot be morally required — that is, it cannot be wrong for me not to make it. My aim in the rest of this essay is to refute this view.

I shall not be arguing directly that contributing to pools for beneficence is morally required. There may be other objections to this claim, for some or maybe even all such pools. But I shall argue that this one fails.

II: Other Treatments

There have been other prominent treatments of the same problem.⁴ Let me preface my own by locating it in relation to them.

Most commonly, the imperceptibilist is dismissed with a counterexample. This is the strategy in Jonathan Glover's seminal discussion: he describes a case in which a large number of conspirators collectively act to harm a large number of victims, but take care to do so in a way which means that none of them individually makes a perceptible difference to any victim.⁵

Subsequent discussion has focused on the question whether we can plausibly oppose the imperceptibilist without arguing for the existence of imperceptible harms and benefits. According to Parfit, we can.⁶ Parfit's own view, like Glover's, is that there can be imperceptible harms and benefits. But he argues that even someone who denies this ought to accept that whether my action is morally required or morally wrong can depend on whether it is, or would be, part of a collective action that perceptibly harms or benefits. His method of argument is to describe imagined examples, make plausible intuitive judgements about them, and use these to derive principles setting out the general conditions under which contributing to collective acts of beneficence is morally required, and contributing to collective acts of harming is wrong. Two of these principles (with Parfit's labels) are:

(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.

(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 people if *enough* of them act in this way, and (4) they would benefit people *most* if they *all* act 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.⁷

These claims are rejected by Michael Otsuka.⁸ Objecting to (C10), he focuses on its clause (4). Unless benefits can be imperceptible, Otsuka argues, it will be possible that withdrawing any single contribution will not reduce the benefit produced by a pool. But perhaps each amount contributed could have produced a perceptible benefit if used in some other way. If so, then given a group some members of which are

contributing to the pool, the benefit produced by that group could always be increased by withdrawing one contribution. However, if everyone contributes, more benefit is produced than if no-one does. Thus the relation "more beneficial than" can generate a cycle when applied to the various outcomes produced by different numbers of contributors, if benefits cannot be imperceptible. If so, there will be no most beneficial outcome that the group could produce.⁹

Otsuka's second objection to Parfit's treatment is that it cannot satisfactorily address what I shall call the "overflow problem". A sensible view will have to avoid the conclusion that I am morally required to add a contribution to a limited-capacity pool which is already full. Parfit's proposal for dealing with this is:

(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.¹⁰

But Otsuka points out that, if withdrawing the contribution of any individual contributor would not reduce the benefits produced by the group, *every* contributor should believe that his joining would make the group too large.

With these objections, Otsuka is attacking the view that the imperceptibilist can be opposed independently of the question whether benefits and harms can be imperceptible. But he goes on to argue that they can be; and that the imperceptibilist's view can therefore be opposed on the grounds that contributions to a collective action of beneficence can themselves produce the most benefit of the actions available to us.¹¹ This leads him to an interesting conclusion. Whether I ought to contribute to a pool for beneficence will depend on the function governing the marginal utility of the benefit being conferred, and on the benefits that might be produced by me if I withhold

my contribution. Suppose (a) the marginal utility of the goods distributed by a limited-capacity pool diminishes in a strictly linear way, and (b) I could give the whole of my contribution in the form of a benefit to someone as badly off as those whom the pool is helping. Under these conditions, I should contribute if I will be contributing to the first half of the pool, but not if I will be contributing to the second half.¹²

Otsuka's appeal to imperceptible benefits thus yields only a limited defence of the moral requirement to contribute to pools for beneficence: he has to deny that we are required to *fill* such pools. Moreover, notice that this defence is limited in a second, much more important respect: it will not extend to pools for beneficence of the actual kind with which we are most familiar. For if Section I was right, my contribution will make *no* difference — not even an imperceptible one — to the beneficiaries of these pools.

In what follows, I argue that Parfit's central claim is right. Even if there are no imperceptible benefits, my action of contributing to a pool for beneficence can still be morally required because it is part of a collective action that confers a perceptible benefit. This thought will extend to cover contributions to pools for beneficence of the actual kind. However, in order to show this, I shall be asking a question which the writers just surveyed leave unasked. They simply take it as obvious that the imperceptibilist is wrong.¹³ The emphasis in their discussions is on saying which is the most plausible form of opposition to the imperceptibilist, and which imperatives of beneficence should be endorsed. But I think these issues can only be settled — and the objections to Parfit avoided — by answering a more fundamental question. Confronted with an imperceptibilist who does not agree that her view is obviously wrong, what can be said to show that it is, beyond the complaint that it leads to verdicts about particular cases that strike "us" (non-imperceptibilists) as false?

III: The Transitivity of Wrongness

In seeking to spell out the case against the imperceptibilist, it is natural to start with the following simple line of thought. Helping needy people is an end which each of us is morally required to pursue: it would be wrong not to do so. But it clearly makes more sense for us to pursue this morally required end cooperatively, as the end of a collective action, for that way it is better achieved. So, if it is wrong not to help, and this is the way in which we ought to be helping, then it is wrong not to help in this way. And that remains true even when this is the only way in which we can help.

This line of thought raises many questions, but not all of them are relevant here. In particular, large questions arise about the scope of its initial claim, that helping needy people is a morally required end. (To which needy people does it apply? What is its relation to other morally required ends? How can the requirement ever be satisfied?) However, for our current purposes, all we need assume is that there are *some* conditions under which helping people is morally required. Obviously, we should need to say what those conditions are if we were seeking to draw conclusions about just which pools for beneficence are the ones to which we are morally required to contribute. However, our aim is the more limited one of refuting the imperceptibilist. And the imperceptibilist's objection is not an objection to the claim that there are conditions under which helping people is morally required, as long as what is in question is *perceptible* help. What we can ask, therefore, is whether, in conditions where its initial claim can be upheld, the argument just sketched can show what is wrong with the imperceptibilist's view.

This argument has a simple three-stage structure. It is wrong not to pursue the end of (perceptibly) helping needy people. But when we can do so through a cooperative pool, that is best. So when we can help through a pool, it is wrong not to do so. A

closer look at this argument, though, reveals that it must be giving us a different set of alternatives at each of these stages. To appreciate this, begin with the initial step. It seeks a starting-point which the imperceptibilist does not contest. But if we are to get an argument for pooling, the claim cannot be that each of us is morally required to give perceptible help, when we could contribute to a pool instead. The argument will have to begin with this claim: if helping individually, without pooling, is our *only* alternative to not helping, then failing to do so would be wrong. From there, the argument moves to observing that if we add the further alternative of helping through a pool, that is better, since it means that we confer a greater perceptible benefit. It then draws a conclusion about the wrongness of not contributing to a pool, even when that is one's only alternative to not helping at all. Thus, the general form of the inference involved here is this:

1. Given only alternatives A and B, it would be wrong to choose A.
2. Given alternatives A, B and C, it would be best to choose C.
3. Therefore, given only alternatives A and C, it would be at least as wrong as before to choose A.¹⁴

(The substitutions we are concerned with are: A = not helping; B = helping individually; C = helping through a pool.)

Are inferences of this form good ones? Under certain conditions, I shall argue, they are. And these conditions do apply to the substitution-instance that concerns us. However, although that gets us part of the way to our response to the imperceptibilist, it will not get us all the way.

Inferences of this general form are not always good: there are some clear counterexamples. Here is one.¹⁵ Suppose that, to fulfil an important promise, I have to

get to the Town Hall by midday, and knowingly doing what will prevent me from keeping this promise would be wrong. And suppose I know that if routes A, B and C are all open, A would be too slow, and C would be quicker than B: thus taking C would be best and it would be wrong to take A. I know as well that if route C were closed, all traffic would be rerouted onto A, which would be slower than ever: taking A would remain wrong, given only alternatives A and B. However, I might also know that, if route B were closed, this would prevent a lot of traffic from getting onto A, and it would now be quick enough, whereas C would now be too slow. Thus, given only alternatives A and C, A would be right and C wrong.

This counterexample relies on the point that restricting my alternatives can alter the consequences of choosing those alternatives, and that this will sometimes alter the goodness or badness of choosing them. Other counterexamples do not rely on this. For instance, it may be the duty of a public official to implement whatever alternative is favoured by a vote; and on many voting systems the result may well violate the suggested inference given above.¹⁶

Thus, there are different ways in which the amount by which three alternative actions A, B and C are better or worse than each other can be altered by the restriction of a choice to two of those alternatives only. This may be altered if the restriction of alternatives alters their consequences; and it may be altered if the evaluative ranking of the alternatives is a function of the preference orderings of the members of a group. We might express what the counterexamples have in common by saying that they exhibit ways in which the "transportability" of evaluative judgements across restrictions of alternatives can be blocked, where:

Evaluative judgements about A, B and C are transportable across restrictions of alternatives if and only if the amount by which any pair of A, B and C are better

or worse than each other is unaffected by whether the third alternative is available or not.

What this gives us, then, is a first assumption that is needed if inferences of the form given above are to be sensible. Evaluative judgements about actions A, B and C must be transportable across restrictions of alternatives.

This deals with the counterexamples discussed so far, which come from a first direction. But it does not handle those that come from a second. These concern supererogatory actions. There are different ways of defining these, but for our purposes the best definition is this: a supererogatory action in a given situation is not morally required, and is morally better than actions which are morally required in that situation. If alternative C is a supererogatory action, then this gives us another way in which inferences of the form given above might be mistaken. Suppose A is doing nothing, B is easily saving a life, and C is giving up my own life to save several: if C is supererogatory, this result will follow.¹⁷ Here, what makes the action supererogatory is its cost to the agent. Clearly, the inference only works if we restrict it to cases that do not involve actions like this. Let us do so, then, by making a second assumption: what makes C better than A and B is not its greater cost to the agent.

Notice that this assumption does not seem to cover all the ways in which an action can be supererogatory, on the definition just given. An action might go beyond what is morally required without being especially costly to me. For example, it seems that being friendly to casual acquaintances is better than not being so, but that not being so, provided it is not a matter of positive rudeness, is not actually morally wrong. It is not as if friendliness involves any extra personal cost, though: it is rather that interpersonal behaviour on this scale, although enough of a moral issue for judgements of what is better and worse to make sense, is not enough of a moral issue to talk sensibly about

wrongness. However, we do not need to expand our assumption to cover this kind of supererogation as well: doing so would be redundant, as we shall see.

To these two assumptions, let me add a third. This is that we are not dealing with cases of moral dilemma, in which whatever you do is wrong. We might write this into the statement of the inference-form by replacing "wrong" with "uniquely wrong", where an action is uniquely wrong, given a set of alternatives, if it is the only action from among those alternatives that is wrong. If an action is *uniquely* wrong, we can infer that it is worse than the alternatives. And if this condition holds, together with the other two that have just been set out, then inferences of the form given above will be good.

To see this, begin with step (1): given only alternatives A and B, it would be wrong to choose A. Invoking our third assumption — that this would not just be wrong, but uniquely wrong — we can infer that it would be worse for me to choose A than B. Next, step (2) tells us that C would be better than either A or B. Given the first assumption, concerning the transportability of evaluative judgements about A, B and C across restrictions of alternatives, we can infer that in the different situation in which my only alternatives are A and C, C would still be better than A.

Now from the fact that I have only two alternatives, and that one of them is better than the other, it does not follow directly that choosing the worse alternative would be wrong. We have seen that there are two reasons for this. What makes the better alternative better might be that it involves a cost to the agent which it would be supererogatory for her to bear; or the difference between the two actions, although morally significant enough to say that one is better, might not be morally significant enough to say that the other is wrong. But the former possibility is ruled out if we invoke our second assumption: that what makes C better than A and B is not its greater cost to the agent. And the latter can also be ruled out, given what we already know. We

saw at step (1) that the difference between A and B is morally significant enough to say that A is wrong, and we know from step (2) that the difference between A and C is even greater (given the transportability assumption). So it could hardly be the case that C is supererogatory in the second way either. If not, we can legitimately move to step (3), inferring that A must remain at least as wrong when the only alternative is C.

Thus inferences of the three-step form are good, provided the three assumptions — concerning transportability, personal cost, and unique wrongness — hold.

IV: Pooling and Transitivity

We have seen the conditions under which transitivity-inferences of the form identified above are good. Now let us consider the following substitutions, in relation to some case of large-scale need:

A = not helping

B = helping individually

C = helping through a pool.

And let us suppose that we are dealing with a case in which we have the relevant substitution-instance of (1):

1'. Given only the alternatives of helping individually and not helping, it would be wrong for us not to help.

(As we have noted, it is a controversial matter for which sorts of need this will be true; but that controversy is not relevant to the debate between us and the imperceptibilist.)

Notice that "helping someone individually" can be taken to span two possibilities:

helping independently, without relying on an agency to administer the help; and helping through an agency, but one which "earmarks" the aid, using my contribution to confer a perceptible benefit on a small number of people.

But now suppose the situation changes: a central pool is set up, together with an organized means of distribution of the resources from that pool. As long as it is more efficient for us to channel our resources through the pool, it will clearly be better for us to do so: for more people will be perceptibly benefited to a greater extent. Even the imperceptibilist, who thinks that only perceptible benefits are morally significant, has to agree. So this gives us the relevant substitution-instance of (2):

2'. Given the alternatives of helping individually, helping through a pool, and not helping, it would be best for us to help through the pool.

But this means that we can infer the corresponding instance of (3), provided the three conditions identified in the previous section hold. And surely they do. First, there seems to be no case for resisting the transportability of evaluative judgements across restrictions of the three alternatives we are considering. The transportability assumption is certainly not violated in either of the ways we examined earlier. The evaluation of these actions does not depend on consequences of them which change when the alternatives are restricted. Nor, obviously, is it a function of the preference-orderings of the members of a group. Unless some further ground for blocking the transportability assumption can be found, it will have to be granted. With the other two assumptions, moreover, we can be more conclusive. What makes helping through the pool better than helping individually is obviously not that it involves greater personal cost: on the contrary, it is crucial that the cost to us of producing the same amount of help is lessened by pooling — we can produce more help with the same amount of effort. And (turning to the third assumption) the case for (1') will be a case for

thinking that not helping is *uniquely* wrong, compared with the alternative of helping individually.

In other words, it is plausible to think that all three conditions are met. And if so, then we can indeed infer:

3'. Given only the alternatives of helping through the pool and not helping, it would be at least as wrong as before not to help.

Often, one's only alternatives are either to help through a pool or not to help at all. But if, had one had the opportunity to help individually, it would have been wrong not to do so; and if it is better for us to help through the pool than to do so individually; then it will be wrong not to help through the pool where this is our only alternative to not helping at all.

V: Individual and Collective Wrongness

Should this argument trouble the imperceptibilist? Not yet. For it invites a clear objection. The imperceptibilist should object that (2') is ambiguous between a collective and an individual reading:

C-2'. If we collectively have the alternatives of helping individually, helping through a pool, and not helping, it would be best for us collectively to help through the pool.

I-2'. If a person has the alternatives of helping individually, contributing to a pool, and doing nothing, it would be best for that person to contribute to the pool.

Now it is true that the imperceptibilist possesses no objection to the collective reading, (C-2'). For, collectively, we could confer a greater perceptible benefit by pooling our

resources than we would by acting independently; so the imperceptibilist has no objection to the plausible claim that, of the alternative actions we could collectively perform, helping through the pool would be best. But what she does object to is the individual reading, (I-2'). The action of any one of us in contributing to the pool would make no perceptible difference to any of the pool's beneficiaries, and so is not the best thing to do. Indeed, she can point out, "helping through a pool", although it may be an option for us collectively, will often not be an option for any of us individually.¹⁸

The problem is that the individual claim is what is needed to derive the corresponding version of claim (3'):

I-3'. If a person's only alternatives are contributing to the pool and doing nothing, it would be wrong to do nothing.

This conclusion cannot be claimed without begging the question, at step (2'), against the imperceptibilist. No objection has been presented to her claim that if my contribution to a collective activity of beneficence confers no perceptible benefit on anyone, it cannot be morally required.

Notice that the imperceptibilist can happily allow that if it were somehow in my power to decide what we should all do, then I ought to decide that we should help through a pool. And she can also allow that, if it were in my power to set up a pool for beneficence or not, but not to change anyone's psychology, then I ought to do so. After all, given that people are not generally imperceptibilists, this would be likely to produce more perceptible benefit. But she should still add that, having set up such a pool, I should not myself be morally required to contribute to it, if this would make no perceptible difference to its beneficiaries.¹⁹ If my own actions would influence many other people to do likewise, then I should contribute. But short of this, there will be nothing wrong in refusing.

We do not yet, therefore, have a good argument against the imperceptibilist. However, we do have one ingredient in that argument. For what we *can* infer from the discussion so far is the collective version of (3[∧]):

C-3[∧]. If we collectively have the alternatives of only either helping through the pool and not helping, it would be wrong for us not to help.

For this conclusion follows from (1[∧]) and (C-2[∧]), which we have seen the imperceptibilist is in no position to reject.

Now this conclusion does not yet seem to be a problem for the imperceptibilist. For it simply encourages her to draw a distinction between the wrongness of individual and collective actions, and to claim that this is a case in which, if we all do nothing, we are collectively acting wrongly, although no individual is acting wrongly.

In support of this claim that individual and collective wrongness can diverge, an appeal might be made to the fact that it seems clear in general that collective transgressions of norms need not entail individual transgressions of those norms. A group can act irrationally even when no individual member of the group is acting irrationally: consider, for instance, cases of intergroup conflict which are mutually disastrous, but where the pressures on individuals to participate in the group activity make this, arguably, the rational thing for each of them to do.²⁰ And more directly, there are plausible examples of situations in which wrongness attaches to the action of a group without distributing onto the actions of its members. Suppose, for example, that we ought collectively to make a gesture of reconciliation to an unjustly treated group. If we do not collectively do so, we are collectively acting wrongly. But it may not be in the power of any member of the group to offer that gesture on his own: what is required is a gesture from the group. And, plausibly, it might be supererogatory for any individual to undertake the political leadership necessary to getting the group to make the gesture.

A lot of further discussion would be needed to make these judgements compelling. But they do at least lend support to the view that the wrongness of a collective action does not always distribute onto the actions of individual group members. I shall accept this in what follows: from the fact that we are collectively acting wrongly, it does not immediately follow that any individual member of the group is acting wrongly.

However, the argument so far does establish that, if there are needy people whom it would be wrong not to help individually, then we are collectively morally required to contribute to pools for beneficence that help them, when our only alternatives are doing this and doing nothing. Given this collective moral requirement, a further argument can be given for thinking that individuals are morally required to contribute to meeting it, and thus that the imperceptibilist is wrong.

VI: Free Riding and Collective Imperatives

From collective imperatives, telling us what we ought collectively to do, individual imperatives to contribute to meeting those collective imperatives do not always follow. However, they often do. One kind of case in which they do so is where failing to contribute would amount to free riding on the production of a collective benefit. Appreciating the relationship between the collective and individual imperatives in this case will show us why, in the different case of pools for beneficence, individual requirements to contribute will follow from collective ones.

Free riding, I shall assume, is a matter of wrongly failing to contribute to a scheme for producing a public good, while enjoying the benefits of the scheme. Just which are the conditions under which failing to contribute to such a scheme is wrong? This is a good, but controversial, question and I shall not discuss it here.²¹ Without going into

this, though, it is not difficult to give a general characterization of what is wrong about paradigm cases of free riding — like abusing an "honour" system for public transport fares by riding without paying. What is wrong about this is that it exhibits a certain kind of unfairness: it involves arrogating to myself a privilege of not paying for benefits for which I rely on others' willingness to pay.

Now if we addressed this complaint to a free rider, he could try replying as follows: "I am not arrogating any special privileges to myself. For I think that each other person ought to be doing what I do. Of course, if everyone did so, then there would be no benefits for anyone. And I agree that it is more sensible for us collectively to act to produce the benefits than to go without them. But I make a distinction between what we collectively ought to do and what each of us individually ought to do. It is good for each of us that enough of us are contributing to what we collectively ought to do. But the fact that this is what we collectively ought to do does not show that anyone ought individually to contribute to our doing so."

However, saying this will clearly not help to address the complaint of unfairness which is to be made against the free rider. The way in which the free rider is arrogating privileges to himself is not that he makes a claim of permission for himself which he is not prepared to make for other individuals. Rather, it is this: he relies on others to do what we ought collectively to be doing, without contributing himself. As he can see, the group to which he belongs ought to be acting to produce the benefit he is enjoying, because the benefit is worth its cost. That collective imperative is being met, but he is leaving the work of meeting it to others. Thus, when he protests that he can accept the collective imperative while consistently denying that this gives rise to any individual imperatives, the response is that this is irrelevant. The complaint against him is not that he is being inconsistent, but that he is violating a requirement of

fairness. Fairness is what tells us to derive the individual imperative from the collective one.

This account of what is wrong in free riding appears to invite an objection. Paradigm cases of free riding do involve benefits which a group ought collectively to be acting to produce. But sometimes, public goods will be produced by a group that ought not to be producing them. Perhaps the benefit is not worth the cost, or there is a moral objection to the means by which it is produced. In these circumstances, just as much as in the more usual ones, it seems that if you take a benefit and then refuse to pay for it, you will still be free riding on the group. An attempt to justify non-payment on the grounds that the group ought not to have produced the good will be beside the point. If you had an objection to the production of the good, you ought not to have taken it. Once you have, failing to pay for it is free riding. This can make it seem that the explanation of what is wrong about free riding cannot after all lie in the unfairness of failing to contribute to the meeting of collective imperatives. For it seems that there can be free riding even when the collective activity in relation to which the free riding occurs is not the meeting of a collective imperative.

However, the wrongness of free riding in this kind of case does actually derive from a collective imperative, but in a less obvious way. Suppose I take some cooperatively produced good, but then refuse to pay for it, on the grounds either that it is not worth the price being demanded, or that it has been produced in a morally objectionable way. If the price was made clear in advance, this will not do: if I had an objection to paying it, I should not have taken the good. For I have violated the convention that taking goods to which a publicly understood price attaches binds one to pay that price. Now notice that this convention is itself something that it clearly does make sense for us collectively to sustain. It confers a benefit on all of us: the second-order benefit of having the opportunity to enjoy first-order goods if one wishes.

It does so by imposing a cost: the cost involved is adherence to the conditional requirement that if one accepts those first-order goods, one must pay the publicly understood price. But notice that now we have identified a convention which *is* one that we ought collectively to observe. The (second-order) benefit it creates is worth its cost, and there is no moral objection to it. But this convention is violated by free riding of the kind we are trying to explain — free riding in relation to a group which ought not collectively to be producing a first-order good. So there is after all the same explanation as before for thinking that the free rider acts unfairly. The free rider *is*, even here, relying on others to do what we ought collectively to be doing, without contributing himself.

What this discussion of free riding shows us, then, is that although collective imperatives do not entail individual ones — there is no inconsistency in recognizing the former but not the latter — there are circumstances in which considerations of fairness allow us to derive individual imperatives from collective ones. And once this is established, the case for thinking that these circumstances include pools for beneficence is straightforward.

In Sections III-V, we saw that the imperceptibilist has no case against thinking that we can be collectively required to help others through a pool. If so, an imperceptibilist who refuses to contribute to the pool while others do so will be relying on others to do what we ought collectively to be doing, without contributing herself. So, in the same general way that the free rider arrogates special privileges to himself, the imperceptibilist is arrogating special privileges to herself. In circumstances in which there is a collective imperative that we ought to be meeting, she is leaving it to others to meet it, without being prepared to do so herself, and without any justification for treating herself differently from them.

Imperceptibilists will want to reply to this as follows: "Whether or not any individual contributes to the pool makes no difference to whether the collective action is performed — the collective imperative is no better met if I contribute than if I do not. In fact, no individual act of participation in the pool is contributing to meeting the collective imperative; so I cannot be relying on anyone's doing this when I refuse to participate in the pool, and I cannot be doing anything wrong in refusing." However, we can concede that no single contribution will have the effect that the collective imperative is better met, while insisting that this is irrelevant to the accusation of unfairness. For this does not show that individual acts of participation in the pool are not contributions to meeting the collective imperative; it only shows that an individual contribution to meeting the collective imperative will not by itself affect what is done by the group. The contribution is perceptible — I put a perceptible amount into the pool — but it has no perceptible effect on what is achieved by the group. Thus there is no good objection to saying that the collective imperative is only being met because others are prepared to make contributions to meeting it: we should simply add that their preparedness is a preparedness to make those contributions even when each has no perceptible effect on what is achieved by the group. Therefore, if I am not myself prepared to make a contribution of this kind, I am relying on others to do what must be done if the collective imperative is to be fulfilled, without being prepared to do so myself. And that is unfair, in the same general way that free riding is unfair.

The imperceptibilist might try to complain that it is wrong to talk about an arrogation of special privileges: she accepts that we are collectively morally required to help through a pool, but denies that *anyone*, and not just herself, is acting wrongly in refusing to contribute to it. But this is no more convincing than the free rider's parallel reply to the same accusation. The objection to the imperceptibilist is not that she is covertly inconsistent — that she is making a claim about what is permissible for her

which she is not prepared to make about every other person. It is that she is unfair. What compels the derivation of individual imperatives from collective ones is not a consideration of consistency, but a consideration of fairness. If that is what we ought to say about the obviously wrong conduct of the free rider, then it equally supports the conclusion that failing to contribute to pools for beneficence can be wrong. Provided we ought collectively to be helping through such a pool, it will be wrong for me to be relying on others to satisfy that imperative, while excepting myself from doing so.

Let me make it clear that I am not claiming that failing to contribute to a pool for beneficence is itself free riding. Free riding arises in relation to collective action the aim of which is to benefit the cooperating group; the aim of a pool for beneficence is not to benefit the cooperating group, but another group. The claim, rather, is that the wrongness of the two different kinds of conduct is explained in the same way. Considerations of fairness dictate deriving individual imperatives from collective ones in the case of free riding; but given this, those same considerations will dictate deriving individual imperatives from collective ones in the different case of pools for beneficence.

VII: The Advantages of this Argument

I have argued that, to understand why contributing to pools for beneficence can be morally required, we should look at the way in which individual imperatives can be derived from collective ones. It is a matter of pulling my weight in what we all ought to be doing. Sections III-V discussed the conditions under which there will be a moral requirement on us collectively to help other people through a pool; Section VI derived the individual imperative from this.

I have not been arguing that individual imperatives can be derived from every collective imperative. In particular, consider the two examples mentioned earlier where it seems arguable that no such derivation can be made: perhaps it is rational for a group to be behaving peacefully, but rational for each group member to be contributing to a conflict; and perhaps as a group it would be wrong for us not to make a gesture of reconciliation, but if we fail to do so no group member acts wrongly. In both of these cases, it is significant that the collective imperative is not being met. If so, no one is relying on others to satisfy that imperative while excepting himself from doing so. Before the collective action is started, none of us is arrogating any special privileges in refusing to get it started. We are collectively acting wrongly, but no individual is taking advantage of others' propensity to contribute to doing what we ought to be doing.²²

Now admittedly, in arguing that the derivation does go through in the case of pools for beneficence, I have simply presupposed the wrongness of free riding. If an imperceptibilist denied this as well, the debate would have to continue. I cannot claim to have been arguing independently of any assumptions about what is obviously wrong. However, it has been shown that the response to the imperceptibilist's view can be better than simply saying that *it* is obviously wrong: this has been given a further justification.

This has led us to the same general conclusion as Parfit: whether my action is morally required can depend on whether it would be part of a collective action that perceptibly benefits people. What remains to be appreciated is the way in which we have avoided Otsuka's two objections.

The first of these concerned the cycle produced by the relation of "more beneficial than" when we consider the outcomes that we could collectively produce. This was a problem for Parfit, since it makes it false that we will produce the most benefit if we

all contribute. However, it is not a problem for the argument given here. This argument does require the claim that, given the options we together have of either collectively helping through a pool or not doing so, the former is best. And *that* is true, even if there cannot be imperceptible benefits. But it does not require the further claim that it is better if all of us contribute than if slightly fewer do. That claim was not needed in order to argue (in Sections III-V) that there are collective imperatives of the following form: we are morally required to help these needy people by pooling our resources. Nor was it needed in order to argue (in Section VI) that it is wrong not to contribute to fulfilling such collective imperatives, while relying on others to do so.

Now it might seem possible to reply to this, in the spirit of Otsuka's argument: "Unless benefits can be imperceptible, it will be true that, whenever I am one of n contributors to the pool, we could collectively produce more benefit if the other $n-1$ contributed while I performed some other perceptibly beneficial action instead. But if this is always collectively better, whatever the value of n , then that will commit you to a collective imperative which recommends it. And from this, the individual imperative to be derived is that I ought not to contribute, whatever the number of contributors." However, my argument relies only on a claim about a collective act-type which is best. It does not claim that, whenever there is a collective action which is better than one we are performing, we ought collectively to perform it instead. That claim would certainly be dubious if the "more beneficial than" relation generated a cycle, as would happen if there are no imperceptible benefits. But that is why it cannot be invoked in a non-question-begging argument against the imperceptibilist, and why I have not done so. It is true that, if there are no imperceptible benefits, there will be no value of n for which it is best that n of us contribute. But when I claim that a collective act-type is best, I am not making a claim of that form. The claim from which I derive the collective

imperative is simply that it is that it is collectively best for us to help by pooling our resources.

Otsuka's second objection to Parfit was the overflow problem. The form in which it can be addressed to me is this. Suppose a limited-capacity pool for beneficence is already full. It can still be the case that, once I add my own contribution, I will be contributing to meeting the collective imperative, even if my doing so has led to an overflow. But then what stops us from saying that if I do not add my contribution, I am relying on others to contribute while excepting myself? It can seem that my argument leads to the faulty conclusion that it is wrong for me not to contribute, even once the pool is full.

However, *this* is not the kind of "contribution" which I am relying on others to make in order for the collective imperative to be fulfilled. For a willingness to add to the pool when that will make it overflow is not what is leading to the discharging of the collective imperative. What I am relying on is a willingness of others to help to fill up the pool before it has been filled. If I am not prepared to make contributions of *this* kind, then there is an argument from fairness for the conclusion that I am acting wrongly. But the same argument will not lead to the conclusion that I am morally required to add to the pool once it is full.

Again, this may seem to invite a reply: "Certainly, once the pool is full, adding to it does not contribute to fulfilling any collective imperative. But, unless we challenge the claim that there are no imperceptible benefits, it will be true that no previous individual contribution to the pool contributes to fulfilling the collective imperative. For any individual contributor, the collective imperative is no better fulfilled if he contributes than if he doesn't."

But this is answered by recalling an earlier point.²³ An individual act of participation may — perceptibly — contribute to meeting the collective imperative

without perceptibly affecting what is achieved by the group. Just how should we describe the contributions to meeting the collective imperative which we are relying on others to make? We are not relying on anyone's making a perceptible difference in the extent to which the collective imperative is being met; for perhaps no one is doing this. But we are relying on others to make small individual contributions to the filling of the pool — and to make these contributions despite the fact that no such contribution will make a perceptible difference to the pool's beneficiaries. It is a willingness of contributors to do this that explains how the collective imperative gets met — and not a willingness to add to the pool once it is full.²⁴

Thus, our argument avoids Otsuka's objections. But, just as significantly, it also avoids the limitations of his conclusion. For one thing, there can be moral requirements on me to make the very last contribution to filling a limited-capacity pool.²⁵ If I am relying on others to do what is needed for us to meet the collective imperative without being prepared to do so myself, then the case for saying that I am acting wrongly is unaffected by how close the pool is to being full. Moreover, we have avoided the second, more important limitation. Our argument can apply to pools for beneficence of all three kinds that we identified at the outset, including those actual ones for which my contribution will make *no* difference — not even an imperceptible one — to their beneficiaries.

In this essay, I have not been arguing directly that contributing to pools for beneficence is morally required. But I have shown that this claim should not be resisted on imperceptibilist grounds. Moreover, we have seen how to settle whether this claim is true, for any particular pool for beneficence. Given a pool for helping a certain group of needy people, it must be determined whether we are collectively required to support it. And Sections III-V showed how to do this: we need to ask whether it would be wrong not to help those people individually, and whether our

helping them through a pool would be better. If so, there is a collective requirement; and if there is a collective requirement, it is wrong for me to stand by while others contribute to meeting it.²⁶

REFERENCES

- Brian Barry, *Political Argument* (London: Routledge & Kegan Paul, 1965).
- Garrett Cullity, "The Life-Saving Analogy", in William Aiken and Hugh LaFollette (eds), *World Hunger and Moral Obligation*, 2nd edition (Englewood Cliffs: Prentice Hall, 1996), pp.51-69.
- Garrett Cullity, "Moral Free Riding", *Philosophy and Public Affairs* 24 (1995), pp.3-34.
- Jonathan Glover, "It Makes No Difference Whether or Not I Do It", *Proceedings of the Aristotelian Society (Suppl.)* 49 (1975), pp.171-90.
- Bart Gruzalski, "Parfit's Impact on Utilitarianism", *Ethics* 96 (1986), pp.760-83.
- Russell Hardin, *One for All: The Logic of Group Conflict* (Princeton: Princeton University Press, 1995).
- Michael Otsuka, "The Paradox of Group Beneficence", *Philosophy and Public Affairs* 20 (1991), pp.132-49.
- Derek Parfit, "Comments", *Ethics* 96 (1986), pp.
- Derek Parfit, *Reasons and Persons* (Oxford: Clarendon Press, 1984).
- Donald H. Regan, *Utilitarianism and Co-operation* (Oxford: Clarendon Press, 1980).
- M.J. Scott-Taggart, "Collective and Individual Responsibility", *Proceedings of the Aristotelian Society (Suppl.)* 49 (1975), pp.191-209.

NOTES

¹ Parfit (1984), p.76.

² No doubt, if we continued increasing these funds, the marginal utility of further increases would eventually become very small; but that is a different point.

³ Glover (1975) emphasizes this kind of case. Note that the three kinds of pool I identify are not mutually exclusive.

⁴ See Barry (1965), pp.328-30; Glover (1975) and Scott-Taggart (1975); Regan (1980), pp.56-65; Parfit (1984), Ch.3; Gruzalski (1986), pp.777-82 and Parfit (1986), pp.846-9; Otsuka (1991).

⁵ Glover (1975), pp.174-5.

⁶ Parfit (1984), Ch.3.

⁷ Parfit (1984), pp.70, 77.

⁸ Otsuka (1991).

⁹ Otsuka (1991), pp.139-41.

¹⁰ Parfit (1984), p.83.

¹¹ Otsuka (1991), pp.145-8.

¹² Otsuka (1991), p.147.

¹³ Otsuka does undertake to argue that harms and benefits may be imperceptible; but he does not argue against the view that only perceptibly beneficent actions can be morally required.

¹⁴ What follows is an improved version of an argument I sketched in Cullity (1996).

¹⁵ I owe it to Howard Sobel.

¹⁶ For the general point, I am indebted to John Broome. For an illustration of it that does not even require intransitivity of preferences, take a constituency of six voters in which 1 votes for the ranking CAB, 2 rank the alternatives ACB, and 3 rank them BCA; and suppose the voting system is regulated as follows:

(1) First preferences are to score 3 points, second preferences 2 points and third preferences 1 point.

(2) Votes giving first preference to an unavailable alternative are to be discounted.

¹⁷ *If C* is supererogatory. I do not think it is always supererogatory for a person to give up his life to save several others. But it does seem sometimes to be.

¹⁸ Although we cannot say this of the third kind of pool for beneficence identified in Section I, we have to say it of the second kind, where contributions have no effect on the pool's beneficiaries. Whether we should say this of the first, dispersing kind of pool depends on whether benefits — instances of helping — can be imperceptible.

¹⁹ She would then be saying that it is best to choose an institution that is only the best because people do not do what is best. But there is no contradiction in this. After all, if people are meanly selfish a certain amount of the time, it will be better to have institutions that take this into account, and better for me, given the choice, to choose such institutions; but this is quite consistent with thinking that I ought never to be meanly selfish myself.

²⁰ For extensive discussion of this phenomenon, see Hardin (1995).

²¹ I do so in Cullity (1995), where I also discuss the relevant sense in which the goods enjoyed by a free rider are "public".

²² This is not to say that whenever a collective imperative is not being met, no individual can be acting wrongly in failing to contribute to meeting it. However, it is clear that if this can be wrong, it does not involve the kind of unfairness that I have been concentrating on. Attention has been drawn here to one way in which individual imperatives can be derived from collective ones. Perhaps there are others.

²³ See Section VI.

²⁴ Once the pool is full, there is nothing wrong with not wastefully adding to it. But, of course, that will not be a justification for not having contributed to filling it in the first place.

²⁵ A further question which needs careful treatment concerns whether I can be required to contribute to a pool for beneficence when I have done so already but others have not. I think that often the answer is Yes, but have not the space to defend that view here.

²⁶ For written comments and discussion, I am grateful to Howard Sobel, John Broome, Derek Parfit and Alex Neill.