THE LIFE-SAVING ANALOGY

Garrett Cullity

"It is morally wrong for affluent people not to help the chronically poor." Many people read this remark in a way that allows them to agree with it, while believing that they're not personally doing anything wrong - perhaps even that no individual is. They read it as maintaining that we are collectively acting wrongly. But on a second reading, it issues a direct personal challenge: It is wrong for affluent *individuals* not to help the world's poor, in failing to contribute their own time and money to voluntary international aid agencies. In this essay, I discuss the most prominent argument for this personal challenge.

To argue that individual non-contribution to aid agencies[1] is wrong, you need to identify a morally significant relation in which each affluent individual stands to destitute people. Your argument can take one of two forms. Let's call it a *derivative* argument if it seeks to derive this relation (one between affluent *individuals* and the destitute) from a morally significant relation in which affluent people *collectively* stand to the destitute. *Non-derivative* arguments, by contrast, seek to establish the existence of a morally significant relation between affluent individuals and the destitute more directly, without mediating the relation through group membership in this way.

One method available to both derivative and non-derivative arguments is to argue by analogy. This method begins by identifying a relation which is clearly morally significant, then argues that, since the relation of the affluent to the destitute is relevantly similar, we should attribute the same moral significance to it. Various analogies are used by different

writers in this way: according to the most forceful of them, non-contribution to aid agencies is like failing to take some simple measure to save a life threatened right in front of you.[2] This is the Life-Saving Analogy.

Arguments which invoke the Life-Saving Analogy to support the wrongness of individual non-contribution I shall call Life-Saving Arguments. These differ, depending on the moral relation which the analogy is being used to establish. Some of them maintain that I *violate the threatened person's right* to be saved, I *infringe a duty* to save him, or both.[3] The best-known Life-Saving Argument, though - Peter Singer's - attributes moral significance to a simpler relation.[4] If you could easily save a stranger's life - you could save a drowning child, say, by wading into a pond - then failing to do so would be seriously wrong. And according to Singer, this shows that the following relation is morally significant: *being able to avert a very great harm to a person* (in this case, the loss of his life) *at small personal cost*. Any plausible normative ethical theory, he maintains, must agree. But every affluent individual stands in this relation towards those destitute people whose lives are threatened by illness or hunger: each of us can help them, at small personal cost, through the agency of aid organizations.[5] Undoubtedly, there are differences between the two cases: we could sum them up by saying that in his example, you could save a life *directly*, whereas with the distant poor you can do so only indirectly. But unless those differences are *morally relevant*, non-contribution to aid agencies remains as seriously wrong as failing to save a life directly.

Now the most fundamental kind of objection to Life-Saving Arguments attacks the method of arguing by analogy itself. This method begins with one kind of situation where it seems "intuitively" (i.e., pretheoretically) obvious that inaction is morally wrong, and then maintains that unless a moral disanalogy can be found, inaction must be regarded as equally wrong in a second kind of situation. But what prevents us from arguing in the reverse direction? It seems intuitively obvious to most people that non-contribution to aid agencies is not seriously morally wrong: why not argue that unless there are morally relevant differences, refusing to save a life directly is not seriously wrong? Alternatively,

why not argue that since it seems intuitively obvious to most people that inaction is seriously wrong in one case but not the other, there *must* be a moral disanalogy? The essence of the objection is this: if you appeal to one set of widespread moral intuitions in order to challenge another, what entitles you to be more confident about the first than the second?

This first, methodological challenge is a serious one, but I have argued elsewhere that it can be met.[6] A different kind of objection accepts the method of argument by analogy itself, but attacks the Life-Saving Analogy in particular. One familiar objection of this kind maintains that aid agencies confer no net benefit on the poor, and perhaps even impose a net cost: contributing to aid agencies, far from being like saving a life, is more like exacerbating the threat.[7] But I shall concentrate here on two less familiar objections. The first questions whether, even if aid agencies do confer a net benefit on the very poor, the effect of my contributions of time and money to those agencies will be the saving of life. And the second questions whether one can begin (as Singer does) with a situation involving a single threat to life, and draw his conclusions concerning a situation involving an enormous number of potential beneficiaries of our help.

Why concentrate on these two objections? One reason is that they challenge more than just Life-Saving Arguments. We shall find that the first objection threatens *any* argument for the wrongness of individual non-contribution, derivative or non-derivative, and whether arguing by analogy or not; and that the issue raised by the second affects all non-derivative arguments too. The other reason is that both objections are good ones: in its simple form, the Life-Saving Argument must be rejected.

I: SAVING LIFE

Singer's argument by analogy, as I read it, maintains that an affluent person's contributions to aid agencies will avert threats to people's lives.[8] However, this claim is

false. The reason for this is not that the net effect of the activity of aid agencies fails to benefit the very poor. (I think a clear case can be mounted for opposing such a view.)[9] Rather, there are two better reasons.

Disaster Prevention and Disaster Relief

The first is simply that most non-government agencies are now primarily concerned with preventing future harm to poor people and improving their lives, rather than conducting emergency relief operations. To prevent the perpetuation of food scarcity, aid agencies have been concerned since the 1960s not simply to provide hungry people with food, but to implement programmes aiming to reduce communities' vulnerability to further crop failures, and to foster improvements in basic medicine and education. Many nongovernment agencies do seek to supplement governmental emergency relief operations, but this is becoming a progressively smaller part of their activities. Most of what they do, therefore, amounts not to providing life-saving aid, but rather preventing the need for it. It is not so much saving a drowning person as funding a swimming education programme.

Clearly, however, this first objection does not apply to the disaster relief activities to which some aid agencies do devote themselves entirely, and many others in part. Here, a second objection applies.

Spreading the Effect of my Donation

Suppose that, in response to a distant food crisis, I donate enough money to an aid agency to sustain one person for its likely duration. What will the effect of my donation be? Hopefully, it will enable the agency to buy more food.[10] But the extra food bought with my money will not be used (nor would it be proper for it to be used) to feed one extra person. It will be sent to a food distribution camp, and shared among the hungry people there. Had I refrained from making my donation, no-one would have failed to

receive food: the available food would have been spread a little more thinly across everyone. And only very slightly more thinly. If there are a thousand people in the camp, their each receiving a thousandth of a food ration more or less each day will not make much difference. Indeed, the effect of this increment of food upon a person's hunger and health is likely to be imperceptible. (Even for those people whose bodies have a fairly definite threshold with regard to malnutrition - so that at a certain level of food intake, reducing it only slightly will put them suddenly in a precarious state - it is unlikely to be my non-contribution which makes this difference, rather than, for instance, the method of food allocation at the camp.)

This is not to deny that contributors to aid agencies collectively make a significant difference to the destitute. But I do not make such a difference. Any hungry person should be quite indifferent to whether I donate or not. Indeed, notwithstanding my far greater wealth, I probably lose more by making such donations than anyone gains from them. Let us call this the imperceptibility objection.[11]

The challenge to the Life-Saving Argument is plain. Its analogy, it seems, should be revised as follows: I'm invited to contribute to a fund for employing lifeguards to patrol a certain stretch of water, but am told that my donation won't make the difference between anyone's living and dying. Surely my refusal seems intuitively less wrong than failing to rescue a person drowning right next to me.

Generalizing the Imperceptibility Objection

The imperceptibility objection, therefore, threatens the Life-Saving Argument. But next, notice that it challenges all other arguments for the same conclusion.

At the outset, I divided arguments for the wrongness of individual non-contribution to aid agencies into two categories. Derivative arguments derive the claim that there is a morally significant relation between affluent individuals and the destitute from the claim

that the affluent collectively stand in a morally significant relation to the destitute; non-derivative arguments do not.

So far, we have been considering one kind of non-derivative argument: an argument by analogy for attributing moral significance to the relation of *being able to avert a very great harm at relatively small cost to yourself*. The imperceptibility objection has shown that the affluent do not stand in this relation even to the starving. Therefore, whether they argue by analogy or not, non-derivative arguments seeking to establish *this* relation fail. But the difficulties can be extended to the other relations which non-derivative arguments seek to establish - a right not to be hungry, for instance, or a duty of beneficence borne by the fortunate towards the needy. If my contributions to an aid agency will not themselves substantially help anyone, how can my not making them violate such a right, or abrogate such a duty?

It is tempting to think that derivative arguments escape these problems. After all, the imperceptibility objection does not impugn our ability *collectively* to avert great harm to the destitute at relatively small cost. The difficulty, though, is to spell out the derivation of *individual* moral wrongness from this or any other collective relation to the destitute. The most natural route for such a derivation would be this: given a group which collectively acts wrongly in failing to achieve a certain result, any member of the group who fails to contribute to achieving that result is individually acting wrongly. But the imperceptibility objection has blocked this route: my donating money to an aid agency will not, in any straightforward sense, contribute to feeding the hungry. If I donate it they will be no less hungry than if I don't.

It might seem that what is needed here is a universalization argument: what makes my individual non-contribution wrong is that if everyone acted in the same way, the destitute would be left unassisted. However, for a start, this pattern of argument seems to run into clear counterexamples. Suppose we ought collectively to build a fence to protect a vulnerable fellow-villager and each of us owns ten fenceposts. If few others are doing anything, surely I'm not morally bound to erect my ten posts anyway? And more

seriously, consider *why* this is a counterexample. What is ridiculous about my unilaterally erecting the posts is that this personally troublesome action makes no contribution to the end of protecting the villager. But as we have seen, there is a clear sense in which my donating money to an aid agency makes no contribution to the end of helping the destitute - it makes no difference to them whether or not I do it - although it makes a difference to me.

The imperceptibility objection, then, challenges *all* arguments for the wrongness of individual non-contribution. However, there is a reply to it.[12]

The Argument from Transitivity.

Consider a famine relief scheme which operates as follows: each donation is used to buy a particular parcel of food, which is allocated to a particular needy individual. No such scheme exists,[13] but if it did, it would clearly circumvent the imperceptibility objection - my contribution would make a perceptible difference to the scheme's beneficiaries. A Life-Saving Argument for the wrongness of failing to contribute to *this* sort of scheme remains unchallenged. The failure to contribute *would* be a failure to avert threats to life.

But there is a good reason why such schemes do not exist. It would be perverse to adopt the earmarking policy in preference to the actual one, of using donations collectively to fund a food supply shared among the occupants of a food distribution camp. The earmarking policy would be perverse (quite apart from considerations of unfairness) because of its inefficiency - which is to say that more people would suffer more greatly under this policy than under the actual one. If so, it is clear what I morally ought to do if I could decide whether aid agencies were to conform to the earmarking model or the existing one. People would be left perceptibly worse off by the earmarking policy; therefore there is a clear reason for anyone concerned about other people's welfare

to choose the non-earmarking policy, and none favouring the alternative. Again, the imperceptibility objection is no objection to this.

But I don't have any such power; so how does this help to address the imperceptibility objection? To see the counterargument, we need to think about the transitivity of judgements of moral wrongness. Consider, first, inferences of the following form:

1. Given only alternatives A and B, it would be uniquely wrong to choose A.

2. Given only alternatives B and C, it would be uniquely wrong to choose B.

3. Therefore given only alternatives A and C, it would be uniquely wrong to choose A.

When I say that doing something would be "uniquely wrong" given a certain set of alternatives, what I mean is that it is the only thing that would be wrong given those alternatives. (That is, if there can be cases where all your alternatives are wrong, this is not one of them.) With this clarification, the above form of inference is easily supported. If, given only two alternatives, choosing one would be uniquely wrong, then it would be worse to choose it; and "worse than" is a paradigm of a transitive relation. If A is worse than B, and B is worse than C, then A must be worse than C.

But now consider a second kind of inference:

1. Given only alternatives A and B, it would be uniquely wrong to choose A.

2. If one had to choose between being given only alternatives A and B, and being given only alternatives A and C, it would be uniquely wrong to choose to be given only alternatives A and B.

3. Therefore given only alternatives A and C, it would be uniquely wrong not to choose A.

Let's apply the same line of thought here. According to premise (2), choosing to be put in a first situation (where one must choose between A and B) would be worse than choosing to be put in a second (where one must choose between A and C). And premise (1) tells us, as before, that choosing A would be worse than choosing B. Suppose for a moment that

choosing A were not worse than choosing C. If so, (2) would tell us that choosing the first situation, where the best one could do is to choose B, would be *worse* than choosing the second, where one could not do better than to choose A. But this implies, given premise (1), that choosing to be able to choose the better of two alternatives would be worse than choosing to be able to choose the worse of the two. And this is false. Therefore the supposition must be false: A must be worse than C. So the conclusion, (3), must be true if the premises are true.

The second form of inference, then, is also valid. And if so, an argument of this form refutes the imperceptibility objection. For A, read "spending all my money on myself"; for B, "contributing some money to an aid agency of the earmarking kind"; and for C, "contributing some money to an aid agency of the existing, non-earmarking kind". We saw earlier that it would be uniquely wrong to keep one's money to oneself, rather than contributing to an earmarking agency, if these were one's only alternatives. But we also saw that it would be uniquely wrong to choose to have earmarking agencies rather than non-earmarking ones, were one given that choice. If so, it must be wrong to keep one's money to oneself instead of contributing to a non-earmarking agency.

This deals with the second, imperceptibility objection to the claim that my contributions to aid agencies save lives. But notice that the same strategy of argument also succeeds against the first - the objection that those aid agencies are primarily working to prevent threats to life instead. For again, given a choice between preventing such threats from arising and averting them once they have arisen, it would be uniquely wrong not to choose the former. However, if averting present poverty-related threats to life were the only option, failing to do so would be uniquely wrong, thanks to the Life-Saving Analogy. The transitivity of moral wrongness again implies that where the preventive option is also available, refusing to take either can hardly be less wrong.

So far, I have considered two reasons for believing that an affluent individual's contributions to aid agencies are unlikely to avert a threat to anyone's life. Many voluntary aid agencies are more concerned with preventing future threats than addressing

present ones, and the effects of my contributions will be spread over many people rather than concentrated on a few. In reply, I have not denied any of this. Non-contribution to aid agencies is not a failure to save life. However, I have shown that non-contribution remains morally *analogous* to the failure to save life: a Life-Saving Argument remains intact. If we consider any contribution large enough to have saved someone's life through an earmarking agency (and it seems that we're talking of something like $30 here),[14] then even if it will not actually save anyone's life, the argument from transitivity shows that failing to make it is as wrong as failing to save life. Any alternative non-derivative or derivative argument can handle the imperceptibility objection in the same way.

II: SAVING LIVES

I have examined a first challenge to the Life-Saving Analogy between non-contribution to aid agencies and the direct and unextenuated failure to save a life. The effect of my contributions of time and money to those agencies will almost certainly not be the saving of life. However, this point has been met, leaving a Life-Saving Argument intact. Now for a second disanalogy. It affects, not *whether* the Life-Saving Analogy supports a conclusion concerning the wrongness of individual non-contribution, but rather the strength of that conclusion.

So far, we have simply been discussing the conclusion that non-contribution to aid agencies is as wrong as letting someone die in front of you. This, I think, is a striking conclusion. However, it is compatible with holding that no-one who has made one donation to an aid agency acts wrongly in stopping there. A Life-Saving Argument which claimed only this would perhaps be criticizing relatively few affluent people.

Singer goes much further. His full conclusion is that "we ought to give until we reach the level of marginal utility - that is, the level at which, by giving more, I would cause as much suffering to myself or my dependants as I would relieve by my gift. This would

mean, of course, that one would reduce oneself to very near the material circumstances of a Bengali refugee."[15] How does he arrive at this extremely demanding claim?

Iteration

His thought is evidently this: the Life-Saving Analogy can be iterated. Suppose that, confronted by two drowning children, I can only be bothered to save one. No doubt, more can be said morally for me than someone who saved neither, but my failure to save the second child remains morally wrong, and for the same reason as before: the cost of helping is trifling compared to what's at stake for him. The Life-Saving Analogy then instructs us to say the same about someone who refuses to make her second contribution to an aid agency. And the same goes for every successive contribution. Not making it will be wrong provided only that, *considered in isolation* (that is, apart from any contributions I have already made), not making it remains comparable to the direct and unextenuated failure to save a life.

It is sometimes maintained that the most that can be morally demanded[16] of an affluent individual towards helping the destitute is that she do her share. "Why must I sacrifice nearly all of my own interests and projects to relieve this suffering when my being required to sacrifice this much is only the result of the indifference of nearly everyone else, and of their failure to help at all[?] This imposes an unfairly large sacrifice and burden on me. What morality requires is a fair distribution, to all those capable of sharing it, of the burden of helping those in need."[17] However, the counterargument is clear.[18] To the example of the two drowning children, add a second bystander. If she walks off and leaves me to deal with both of them, then no doubt she behaves contemptibly. However, this surely does not allow me to save the first child and abandon the second, in accordance with my share of the required help. The underlying point is simply that, when an accusation of immoral callousness is made against someone who won't avert threats to other people's lives, protestations concerning his share of the cost

are irrelevant. Given the Life-Saving Analogy, the corresponding claim concerning world poverty - that morality demands only doing one's share towards alleviating it, even when one knows that others are not doing theirs - is no more credible.

Now there is one way in which, even on the iterative reading of the Life-Saving Argument, whether further contributions to aid agencies are morally demanded of me will depend on my previous contributions. The greater my contributions to aid agencies, the poorer I'll become, and as I become poorer, donating the same amount of money will become a progressively greater sacrifice. Eventually, that sacrifice might become too severe to demand it of me even to save a life directly. How severe is that? The common view seems to be that if saving someone's life would mean endangering my own, suffering a serious and permanent physical injury, or something comparable, and *this* deters me, then my inaction would not be wrong - rather, saving the life would have been heroic. If this is right, then once the sacrifice in making a further $30 donation has become *this* severe, my subsequent non-contribution will no longer be wrong. On the iterative reading, then, I will be morally required to keep giving money until either I reach this point, or there are no more destitute people to help. And the latter condition is unlikely to be met soon.[19]

Thus, on this iterative interpretation of the Life-Saving Argument, its full conclusion is the following:

The Iterative Conclusion

Ceasing to contribute to aid agencies will only be permissible when I have become so poor that making any further $30 contribution (one large enough to be comparable to averting a threat to someone's life) would itself be such a great sacrifice that it would not be wrong to let someone die in front of me at that cost.

The conventional view of how great that sacrifice must be has just been described. To be sure, this view has been attacked as unjustifiably lenient - Singer evidently thinks so, in concluding that we ought to reduce ourselves to the level below which we would cause more suffering to ourselves than we would relieve in our beneficiaries.[20] But notice that,

even on the conventional view of the magnitude of a heroic sacrifice, the Iterative Conclusion generates severe demands. For even on the conventional view, it demands that I reduce myself to the level below which giving up another $30 would be like submitting myself to a serious and permanent physical injury. And surely I can only claim to have reached this point by becoming very poor indeed.

Aggregation

We wanted to know when the Life-Saving Argument will allow us to *stop* contributing to aid agencies. The Iterative Conclusion seems a persuasive answer. Its persuasiveness comes from its simplicity: if non-contribution is morally comparable to not saving a life directly, and if the grounds for the wrongness of not saving a life are iterative, then the grounds for the wrongness of non-contribution must also be iterative.

But isn't there a way of answering our question which is simpler still? According to the Life-Saving Argument, non-contribution is morally comparable to not saving lives directly; we want to know when we may permissibly stop contributing; so why not simply ask: when may we permissibly stop saving lives directly? And putting the matter this way reveals that the Life-Saving Analogy with which we have been working can be improved. A potential contributor to aid agencies is more closely analogous to someone confronted by a great many drowning people than to someone confronted by one. So why not simply examine our intuitions concerning the more closely analogous case? Suppose, then, that what I come across is not a pond containing one or two drowning children, but the nightmarish scene of a lake, or even a sea, teeming with them. And suppose (to complete the analogy) that many other people could help to save them, but relatively few are doing so. But now, it is surely far from intuitively obvious that it would be wrong of me not to spend practically every waking moment saving lives. This is what the iterative approach would require: I may only stop when the five minutes necessary to save one further individual would itself be a severe sacrifice. But if there is any intuitive response to such

a case, it is surely not this. No doubt, saving no one would be wrong, but would it be obviously wrong, say, to spend my mornings pulling people out of the water and my afternoons pursuing my own life?

It might seem that if intuition says this, intuition must be confused. If we think that the case against abandoning the second of two drowning children is the same as the case against abandoning the first, how can we resist the demanding conclusion concerning the case where many lives are threatened? If the moral considerations relating to a second victim simply iterate those relating to the first, they must be iterated for every subsequent victim as well.

However, there need be no confusion. For we needn't explain the wrongness of abandoning the second of two drowning children in terms of iteration. The iterative approach uses the following method to derive a conclusion concerning what an agent is morally required to do when presented with more than one person in extreme need: begin with what she would have been required to do for a single needy person, then iterate this for every other. But there is an alternative approach to such situations: begin instead by assessing the magnitude of the overall collective need of the people who could be helped, then ask directly what overall sacrifice can be morally demanded of the agent in response to that collective need. This aggregative approach, to be sure, does not as such preclude conclusions as strong as or even stronger than those generated by the iterative approach. But versions of it do yield weaker conclusions; and we saw some intuitive support for one of them in considering the imagined case of a sea of drowning children. According to such aggregative approaches, an agent may permissibly stop saving lives on the ground that she has given up as much as is demanded of her on their collective behalf, even though the sacrifice involved in saving one more life remains small. And on such an approach, the judgements about the two drowning children and the sea of drowning children can be reconciled. The complaint against someone who saves only one of two drowning children will not concern the trifling cost of saving the second child, but the trifling cost of saving *both* lives. This removes the obstacle to agreeing that the failure to

save the second child is wrong, while rejecting the extreme demands generated by the iterative approach. The two intuitions are not contradictory after all.

The aggregative approach draws a different conclusion from the Life-Saving Argument. We can state it thus:

The Aggregative Conclusion

Ceasing to contribute to aid agencies will only be permissible when I have become so poor that any further contribution would make my *total* sacrifice greater than can be demanded of me to save other people's lives.

This formulation leaves it open just how the magnitude of the sacrifice which can be demanded of me relates to the number of lives to be saved. According to the very weakest version of the aggregative approach, this magnitude is constant, irrespective of the number of lives in question: I'm not required to sacrifice more for a hundred people than I am for one. However, only the weakest version claims this. There is a range of aggregative approaches which are progressively more demanding than this weakest version, but whose conclusions remain weaker than the Iterative Conclusion.

Generalizing the Issue between the Two

This gives us the Life-Saving Analogy's second challenge. A potential contributor to an aid agency is more like someone confronted by many drowning people than someone confronted by one, and this appears to support a weaker, Aggregative Conclusion rather than the stronger, Iterative Conclusion sought by Singer. In discussing the first challenge - that my contributions to aid agencies will not save life - I showed that it threatened not only the Life-Saving Argument, but all other arguments for the wrongness of individual non-contribution as well. Does the issue between the Aggregative and Iterative Conclusions also generalize to those other arguments, affecting this time not their validity as such, but the strength of their conclusions?

Yes: but this time, not all of them. Non-derivative arguments do all seem to confront this issue. Recall that they assert a morally significant relation between an affluent individual and the destitute which makes her non-contribution to aid agencies wrong, and which is not derived from a morally significant relation between the affluent collectively and the destitute. Now either the asserted relation holds between the affluent individual and destitute individuals, or it holds between her and the destitute taken collectively. If the former, then it is hard to see how the moral relation - whether it concerns rights, duties, or potential benefaction - can fail to be the same in respect of *each* destitute person. But then she stands to each of them in a relation whose moral significance suffices to make her non-contribution to aid agencies wrong; if so, the route to an extremely demanding conclusion by iteration lies open. If the latter, then it is hard to see how my making a contribution which is morally comparable to saving only one life can fulfil any duty I owe to the destitute collectively, satisfy their collective right to be helped, or meet the demand incurred through any other moral relation I bear to them collectively; again, the case for iteration seems clear.

With derivative arguments, however, the situation is different. They derive the wrongness of my non-contribution from my membership of an affluent group which bears a collective responsibility to help the destitute. It is easy to see how the derivative individual demand could simply be to do my share towards the group's discharging that responsibility.

(However, we have seen the implausibility of requiring me only to do my share if a second bystander leaves me to deal with two drowning children - which is to say that it's implausible to derive the moral demand on me in this case from a collective responsibility. If the Life-Saving Analogy can be upheld, as I claim, then the attempt to derive the wrongness of non-contribution from collective responsibility is equally implausible.)

Adjudicating between the Two

I have suggested that intuition favours the Aggregative over the Iterative Conclusion. When considering someone confronted directly by many threatened lives, we do not readily draw the severely demanding analogue of the Iterative Conclusion. But how strong a source of support for the Aggregative Conclusion is this intuitive judgement?

It is certainly a serious blow to any argument for the Iterative Conclusion which uses Singer's methodology. As we noted at the outset, he begins with a case where the moral status of inaction seems intuitively obvious, then argues that unless a morally relevant disanalogy can be established, we should accord the same moral status to non-contribution to aid agencies. This methodology would seem committed to following our intuitive judgements about the more closely analogous case, and drawing the weaker conclusion. Of course, nothing privileges these intuitive judgements - especially when they concern such a quickly described and far-fetched situation. However, as I have already suggested, the Life-Saving Analogy itself cannot support such an attack. One may try arguing that since not saving a first person at small cost is wrong, and there is no good reason against saying the same for any subsequent threatened person, the weaker conclusion must be mistaken; but this begs the question. According to the aggregative approach, the point about a lone threatened person is not that *he* can be saved at small cost, but that the entire number of threatened people can be.

Thus, the Aggregative Conclusion is not refuted by the Life-Saving Argument itself. There is, however, a strong and simple challenge to it, of the following form. Why should my failure to save the hundred-and-first, or the thousand-and-first person be any more excusable than my failure to save the first, if the cost remains trifling compared to what is at stake for him? My having already saved a thousand lives does nothing to alter the fact that I could save *this* person at an insignificant cost. Intuition is telling me, in effect, to lump all the threatened people together, and assess my sacrifice in relation to this collective entity; but this ignores the very real plight of the individuals conglomerated in this way.

This is hardly the end of the story. It is easy to think of replies offered by various normative ethical theories. An indirect consequentialist may appeal to the beneficial consequences of setting ourselves achievable moral standards.[21] Theories of a contractualist kind may support a resuscitation of the requirement that each of us does his moral share.[22] And some virtue theorists will want to argue as follows: my having already saved a thousand lives may not alter the fact that the person in front of me could be saved at an insignificant cost, but it does help to show that I am not a callous person, and therefore (since the moral assessment of actions derives from that of agents) that my leaving the thousand-and-first person is not wrong. What is difficult, though, is to see how to resolve the issue between the Aggregative and Iterative Conclusions without examining the plausibility of such general normative theories.[23] So I'll have to leave it unresolved here.

Conclusion

I have concentrated on one kind of argument for the wrongness of affluent individuals' not contributing money to voluntary international aid agencies: one which draws an analogy between failing to contribute money to aid agencies and failing to make some small effort to save a life threatened right in front of you. However, in examining two different attacks on this analogy, I have shown that both raise challenges of wider significance. The first attack established that the effect of my contributions to aid agencies will almost certainly not be the saving of anyone's life. This threatened to undermine *all* arguments for the wrongness of individual non-contribution. However, my reply, appealing to relations of transitivity between our moral judgements, left a Life-Saving Argument intact.

The second attack opposed the attempt to use an analogy concerning a threat to a single life to support an extremely demanding conclusion concerning threats to many lives. This challenge concerned not the validity of the argument as such, but the strength

of its conclusion; again, we found the issues here to be of wider significance, applying to all non-derivative arguments.

Whether the second attack succeeds is unclear. Even if it does, though, the Life-Saving Analogy still supports the Aggregative Conclusion. At its very weakest, this claims that you are morally required to keep contributing your time and money to aid agencies, until you have become so poor that any further contribution would make your total sacrifice so great that it would not be wrong to let someone die in front of you at that cost. This is still a much stronger standard than most of us are prepared to live up to.

REFERENCES

BARRY, Brian, 1982. "Humanity and Justice in Global Perspective", *Nomos* 24: 219-52.

BENNETT, Jonathan, 1981. "Morality and Consequences", in S.M. McMurrin (ed.), *The Tanner Lectures on Human Values* 2 (Cambridge, Cambridge University Press): 45-116.

BROCK, Dan W., 1991. "Defending Moral Options", *Philosophy and Phenomenological Research* 51: 909-913.

BROWN, Peter G., 1977. "Food as National Property", in Brown and Shue (1977): 65-78.

BROWN, Peter G., and SHUE, Henry, 1977. Eds, *Food Policy: The Responsibility of the United States in the Life and Death Choices* (New York, The Free Press).

CULLITY, Garrett, 1994. "International Aid and the Scope of Kindness", *Ethics* 105, pp.99-127.

DRÈZE, Jean and SEN, Amartya, 1989. *Hunger and Public Action* (Oxford, Clarendon Press).

DRÈZE and SEN, 1990. Eds, *The Political Economy of Hunger, Vol.2: Famine Prevention* (Oxford, Clarendon Press).

EHRLICH, Paul, 1971. *The Population Bomb* (New York, Ballantine Books).

FAO, 1994. Food and Agriculture Organization of the United Nations, *Commodity Review and Outlook 1993-94* (Rome, FAO).

FEINBERG, Joel, 1970. *Doing and Deserving* (Princeton, Princeton University Press).

FISHKIN, James S., 1982. *The Limits of Obligation*, (New Haven, Yale University Press).

GLOVER, Jonathan, 1975. "It Makes No Difference Whether or Not I Do It", *Proceedings of the Aristotelian Society* (Suppl.) 49: 171-90.

GOODIN, Robert E., 1985. *Protecting the Vulnerable: A Reanalysis of Our Social Responsibilities* (Chicago, University of Chicago Press).

GOROVITZ, Samuel, 1977. "Bigotry, Loyalty, and Malnutrition", in Brown and Shue (1977): 129-42.

HARDIN, Garrett, 1974. "Lifeboat Ethics: The Case Against Helping the Poor", *Psychology Today*; repr. in Rachels (1979b: 279-91).

HARDIN, 1983. "Living on a Lifeboat", in Narveson (1983: 166-78).

HOOKER, Brad, 1991. "Rule-Consequentialism and Demandingness: A Reply to Carson", *Mind* 100: 270-76.

KAGAN, 1989. *The Limits of Morality* (Oxford, Clarendon Press).

KAGAN, 1991. "Replies to My Critics", *Philosophy and Phenomenological Research* 51: 919-28.

LUCAS, George R., Jr., and OGLETREE, Thomas W., 1976. Eds, *Lifeboat Ethics: The Moral Dilemmas of World Hunger* (New York, Harper & Row).

MASKREY, Andrew, 1989. *Disaster Mitigation: A Community-Based Approach* (Oxford, Oxfam Publications).

MEADOWS, Donella H., MEADOWS, Dennis L., RANDERS, Jorgen, and BEHRENS, William W., III, 1972. *The Limits of Growth: A Report for the Club of Rome's Project on the Predicament of Mankind* (London, Pan Books).

MEIER, Gerald M., 1984. *Emerging from Poverty* (New York, Oxford University Press).

MURPHY, Liam B., 1993. "The Demands of Beneficence", *Philosophy and Public Affairs* 22: 267-92.

NAGEL, Thomas, 1977. "Poverty and Food: Why Charity is Not Enough", in Brown and Shue (1977): 54-62.

O'NEILL, Onora, 1975. "Lifeboat Earth", *Philosophy and Public Affairs* 4: 273-92.

PADDOCK, Paul, and PADDOCK, William, 1967. *Famine - 1975!* (Boston, Little, Brown & Co.).

PARFIT, Derek, 1984. *Reasons and Persons* (Oxford, Clarendon Press).

RACHELS, James, 1979. "Killing and Starving to Death", *Philosophy* 54: 159-71.

SEN, Amartya, 1981. *Poverty and Famines: An Essay on Entitlement and Deprivation* (Oxford, Clarendon Press).

SINGER, H., WOOD, J., and JENNINGS, T., 1987. *Food Aid: The Challenge and the Opportunity* (Oxford, Clarendon Press).

SINGER, Peter, 1972. "Famine, Affluence and Morality", *Philosophy and Public Affairs* 1: 229-43.

WHELAN, John M. Jr., 1991. "Famine and Charity," *The Southern Journal of Philosophy* 29: 149-166.

THE WORLD BANK, 1990. *World Development Report 1990: Poverty* (New York, Oxford University Press).

NOTES

[1] I use "aid agencies" throughout to refer to voluntary, non-government, international aid agencies. "Non-contribution" will always refer to affluent individuals' failure to contribute time and money to aid agencies of this kind.

[2] For other arguments by analogy, see Nagel (1977: 58) and Brown (1977: 71), where the analogy concerns property-ownership, and Gorovitz (1977: 132-3, 141), where it concerns bigotry.

[3] See e.g. O'Neill (1975).

[4] Singer (1972). For a similar argument for an even stronger conclusion, see Rachels (1979).

[5] At the time of writing, the last year for which Oxfam was able to supply me with an estimate of the cost of sustaining the lives of its beneficiaries was 1991. In that year, it claimed that £23 would feed one person for six months in Ethiopia, Mozambique, Sudan, Angola, Malawi and Liberia - that is, about a dollar-and-a-half per week.

[6] See Cullity (1994).

[7] See e.g. Hardin (1974) and (1983), Lucas and Ogletree (1976) (especially the contributions by Fletcher, Englehardt and Hardin), Paddock and Paddock (1967), Ehrlich (1971), and Meadows *et.al*. (1972).

[8] Another reading is possible. Singer's conclusion is clearly that *individual* non-contribution is wrong, but he might be read as offering a derivative argument for it - as inferring it from our *collective* ability to avert very great harm to the destitute at a relatively insignificant cost. However, the difficulties for the inference from collective to individual wrongness are discussed below.

[9] The extensive literature supporting this case includes Sen (1981), Drèze and Sen (1989) and (1990), Meier (1984), Singer *et al.* (1987) and Maskrey (1989).

[10] I say "hopefully" because a donation of $20-$30 is actually rather unlikely to do this. World food trade is conducted in tonnes, and at 1993 prices of $123, $137 and $250 for a tonne of wheat, cassava or rice, the likelihood that *my* donation will make the difference between an aid agency's buying an extra tonne and not doing so is small. (FAO 1994: 76, 85, 74.) However, this point by itself does not undermine the Life-Saving Analogy: if I know that spending a small amount of money has a 10% chance of saving several lives, it still seems wrong to refuse.

[11] For versions of this objection, see Whelan (1991: 158-61), and Goodin (1985: 162-3). However, neither concludes that contributing nothing to aid agencies is morally acceptable, and it is unclear why not. Why should non-contribution be wrong if it makes no difference to anyone whether or not I contribute?

[12] For a different response to the imperceptibility objection, see Glover (1975) and Parfit (1984: 75-86). They try to show that it has consequences which even its proponents will find counterintuitive.

[13] Child-sponsorship schemes for the relief of poverty do - although most of these target aid at an entire community rather than a single child or family.

[14] See note [5] above.

[15] Singer (1972: 241).

[16] What I mean in saying that an action is morally demanded is simply that it would be morally wrong not to perform it.

[17] Brock (1991: 912), who calls this the "Why me?" objection. See Kagan (1991: 924-5) for a reply.

[18] This observation is not new. See Feinberg (1970: 244); Bennett (1981: 84); Fishkin (1982: Ch.10); Barry (1982: 222); Goodin (1985: 134-44).

[19] The World Bank estimated in 1990 that over a billion people had an annual income of less than $370, and that 630 million had less than $275 (in 1985 "purchasing parity power" U.S. dollars). The $275-370 range spans its estimated poverty lines - the incomes required to sustain a minimum standard of nutrition and other basic necessities, as well as the estimated cost of "participating in the everyday life of society" - for a number of countries among those with the lowest average incomes: Bangladesh, Egypt, India, Indonesia, Kenya, Morocco and Tanzania. See The World Bank (1990: 26-7).

[20] See Singer (1972: 241). For a detailed challenge to the conventional view concerning heroic sacrifices, see Kagan (1989).

[21] See Hooker (1991).

[22] See Murphy (1993: 290-2), who seeks to distinguish affluent people's relation to the destitute from that of two bystanders to two drowning people.

[23] Difficult, but not impossible, it seems to me. Defending a non-theoretical resolution of the issue between the Aggregative and Iterative Conclusions requires much more space than I have here, though.