8

Rich and Poor

SOME FACTS ABOUT POVERTY

At the end of the twentieth century, the World Bank sent out a team of researchers to record the views of 60,000 women and men living in extreme poverty. Visiting seventy-three countries, the research team heard, over and over, that poverty meant these things:

- You are short of food for all or part of the year, often eating only one meal per day, sometimes having to choose between stilling your child's hunger or your own, and sometimes being able to do neither.
- You can't save money. If a family member falls ill and you need money to see a doctor, or if the crop fails and you have nothing to eat, you have to borrow from a local moneylender; he will charge you so much interest that the debt continues to mount, and you may never be free of it.
- You can't afford to send your children to school; or if they do start school, you have to take them out again if the harvest is poor.
- You live in an unstable house, made with mud or thatch that you need to rebuild every two or three years, or after severe weather.
- You have no close source of safe drinking water. You have to carry it a long way, and even then, it can make you ill unless you boil it.

Along with these material deprivations goes, very often, a humiliating state of powerlessness, vulnerability and a deep sense of shame or failure.

Extreme poverty, as defined by the World Bank, means not having enough income to meet the most basic human needs for adequate food, water, shelter, clothing, sanitation, health care or education. In 2008, the

Bank calculated that this requires a daily income that is the purchasing power equivalent of about US$1.25 per day in the United States. This is not the foreign exchange equivalent of US$1.25, which might not be so bad, because as everyone who travels from a rich country to a poor one knows, the currencies of rich countries often have much greater purchasing power in poor countries. The World Bank's definition takes that difference into account: the poor earn only as much as will buy, in their currency, the quantity of necessities that $1.25 will buy in the United States. The bank estimates that 1.4 billion people have less income than this.

In industrialized countries, people are poor by comparison to others in their society. Their poverty is relative – they have enough to meet their basic needs and usually access to free health care as well. The 1.4 billion people living in extreme poverty in developing countries are poor by an absolute standard: they have difficulty meeting their basic needs. Absolute poverty kills. According to UNICEF, the United Nations International Children's Emergency Fund, 8.8 million children under five years old died from avoidable, poverty-related causes in 2008. That comes to 24,000 – think of it as a football stadium full of children – dying unnecessarily every day. (The number of children dying has been falling steadily since the 1960s but still remains far too high.) Millions of adults also die because of absolute poverty. Life expectancy in the rich nations is now seventy-eight years; in developing countries, it is around fifty. When absolute poverty does not cause death, it still causes misery of a kind not often seen in the affluent nations. Malnutrition in young children stunts both physical and mental development. Millions of people on poor diets suffer from deficiency diseases, like goiter, or blindness caused by a lack of vitamin A. The food value of what the poor eat is further reduced by parasites such as hookworm and ringworm, which are endemic in conditions of poor sanitation and health education.

Death and disease apart, absolute poverty remains a miserable condition of life, with inadequate food, shelter, clothing, sanitation, health services and education. This is the 'normal' situation of our world. At least ten times as many people died from preventable, poverty-related diseases on September 11, 2001, as died in the terrorist attacks on the World Trade Center and the Pentagon on that black day. The terrorist attacks led to trillions of dollars being spent on the 'war on terrorism' and on security measures that have inconvenienced every air traveller since then. The deaths caused by poverty were ignored. So whereas very few people have died from terrorism since September 11, 2001, approximately 30,000 people died from poverty-related causes on September 12, 2001, and

on every day between then and now, and will die tomorrow. Even when we consider larger events, like the Asian tsunami of 2004, which killed approximately 230,000 people, or the 2010 earthquake in Haiti that killed up to 200,000, we are still talking about numbers that represent just one week's toll for preventable, poverty-related deaths – and that happens fifty-two weeks in every year.

SOME FACTS ABOUT AFFLUENCE

We can juxtapose a picture of 'absolute affluence' against this picture of absolute poverty. Those who are absolutely affluent are not necessarily affluent by comparison with their neighbours, but they have more income than they need to provide themselves adequately with all the basic necessities of life. After buying (either directly or through their taxes) food, shelter, clothing, basic health services and education, the absolutely affluent still have money to spend on luxuries. The absolutely affluent choose their food for the pleasures of the palate, not to stop hunger; they buy new clothes to look good, not to keep warm; they move house to be in a better neighbourhood or have more space for the children to play, not to keep out the rain; and after all this, there is still money to spend on home entertainment centres and exotic holidays.

At this stage, I am making no ethical judgments about absolute affluence; I am merely pointing out that it exists. Its defining characteristic is a significant amount of income above the level necessary to provide for the basic human needs of oneself and one's dependents. By this standard, the majority of citizens of Europe, North America, Japan, Australia, New Zealand and the oil-rich Middle Eastern states are all absolutely affluent. There are also hundreds of millions of affluent people in countries like China, India and Brazil, although there is also extreme poverty in those countries. These affluent people have wealth that they could, without threatening their own basic welfare, transfer to the extremely poor.

At present, very little is being transferred. In 1970, the United Nations General Assembly set a modest target for the amount of foreign aid that the rich nations should give: 0.7 percent of Gross National Income, or 70 cents for every hundred dollars a nation earns. Forty years later, only Denmark, Luxembourg, The Netherlands, Norway and Sweden have reached that level. In 2008, the United States and Japan, the two largest economies among the affluent nations, gave only 0.19 percent, or 19 cents in every $100 they earned. Australia and Canada did only slightly better, at 0.32 percent, whereas France, Germany and Britain were around the

average for affluent nations, giving between 0.38 and 0.43 percent. In comparison to their income, what the rich nations are giving is relatively trivial.

THE MORAL EQUIVALENT OF MURDER?

These facts suggest that, by giving far less than they could, rich people are allowing more than a billion people to continue to live in conditions of deprivation and to die prematurely. This conclusion applies not only to governments but to each affluent individual, for each of us has the opportunity to do something about the situation; for instance, to give our time or money to voluntary organizations that are helping to provide health care, safe drinking water, education and better agricultural techniques for the poor. If, then, allowing someone to die is not intrinsically different from killing someone, it would seem that we are all murderers.

Is this verdict too harsh? Many will reject it as self-evidently absurd. They would sooner take it as showing that allowing to die cannot be equivalent to killing than as showing that living in an affluent style without contributing to an aid agency is ethically equivalent to going over to Ethiopia and shooting a few peasants. They point to several significant differences between spending money on luxuries, when we could use it to save lives, and deliberately shooting people. Let us look at some of these differences and then consider which of them really are morally significant.

First, the motivation will normally be different. Those who deliberately shoot others go out of their way to kill. Motivated by malice, sadism or some equally unpleasant motives, they want their victims dead. A person who buys an iPod presumably wants to enhance her enjoyment of music – not in itself a terrible thing. At worst, spending money on luxuries instead of giving it away indicates selfishness and indifference to the sufferings of others, characteristics that may be undesirable but are not comparable to actual malice or similar motives.

Second, it is not difficult for most of us to act in accordance with a rule against killing people: it is, on the other hand, very difficult to obey a rule that commands us to save all the lives we can. To live a comfortable or even luxurious life, it is not necessary to kill anyone, but we do have to allow to die some whom we might have saved, for the money that we need to live comfortably could have been given away. Thus, the duty to avoid killing is much easier to discharge completely than the duty to save. Saving every life we possibly could save would mean cutting our

standard of living down to the bare essentials needed to keep us alive.* To discharge this duty completely would require a degree of moral heroism utterly different from that required by mere avoidance of killing.

A third difference is the greater certainty of the outcome of shooting when compared with not giving aid. If I point a loaded gun at someone at close range and pull the trigger, it is virtually certain that the person will be killed; whereas the money that I could give might be spent on a project that turns out to be unsuccessful and helps no one.

Fourth, when people are shot, there are identifiable individuals who have been harmed. We can point to them and to their grieving families. When I buy my iPod, I cannot know who my money would have saved if I had given it away.

Fifth, it might be said that the plight of the hungry is not my doing, and so I cannot be held responsible for it. The starving would have been starving if I had never existed. If I kill, however, I am responsible for my victims' deaths, for those people would not have died if I had not killed them.

These differences need not shake our previous conclusion that there is no intrinsic difference between killing and allowing to die. They are extrinsic differences, that is, differences normally but not necessarily associated with the distinction between killing and allowing to die. We can imagine cases in which someone allows another to die for malicious or sadistic reasons. We can imagine a world in which there are so few people needing assistance and they are so easy to assist, that our duty not to allow people to die is as easily discharged as our duty not to kill. We can imagine situations in which the outcome of not helping is as sure as shooting. We can imagine cases in which we can identify the person we allow to die. We can even imagine a case of allowing to die in which, if I had not existed, the person would not have died – for instance, a case in which if I had not been in a position to help (though I didn't help), someone else would have been in my position and would have helped. These imaginary situations aside, however, it is true that the extrinsic differences that *normally* mark off killing and allowing to die help to explain why we *normally* regard killing as much worse than allowing to die. To explain our conventional ethical attitudes, however, is not to justify

* Strictly, we would need to cut down to the minimum level compatible with earning the income which, after providing for our needs, left us most to give away. Thus, if my present position pays me $100,000 a year but requires me to spend $30,000 a year on living in a more expensive location than I otherwise might, I cannot save more people by moving to an inexpensive rural area if that will mean taking a job that pays only $60,000.

them. Do the five differences not only explain, but also justify, our attitudes? Let us consider them one by one.

(1) Take the lack of an identifiable victim first. Research has shown that people offered an opportunity to give to a poor child are more likely to give if they are shown a photograph of the child and told her name and age than if they are not given any identifying details. But this may show no more than that, during the millions of years in which our ancestors lived in small face-to-face groups, we developed an instinctive response to help individuals. In contrast, we did not develop any response to giving more anonymous forms of aid, for which there was no opportunity anyway. Should this make any difference to our ethical obligations? Suppose that I am a travelling salesperson, selling tinned food, and I learn that a batch of tins contains a contaminant, the known effect of which, when consumed, is to double the risk that the consumer will die from stomach cancer. Suppose I continue to sell the tins. My decision may have no identifiable victims. Some of those who eat the food will die from stomach cancer. The proportion of consumers dying in this way will be twice that of the community at large, but who among the consumers died because they ate what I sold, and who would have contracted the disease anyway? It is impossible to tell; but surely this impossibility makes my decision no less reprehensible than it would have been had the contaminant had more readily detectable, though equally fatal, effects. Moreover, if this is true for killing an unidentifiable individual, why should it be any different for failing to save one?

(2) The lack of certainty that by giving money I could save a life does reduce the wrongness of not giving, by comparison with deliberate killing; but it is insufficient to show that not giving is acceptable conduct. The motorist who speeds through pedestrian crossings, heedless of anyone who might be on them, is not a murderer. She may never actually hit a pedestrian; yet if she knowingly risks killing an innocent person, what she does is very wrong indeed.

(3) The idea that we are responsible for our acts but not for our omissions is more puzzling. On the one hand, we feel ourselves to be under a greater obligation to help those whose misfortunes we have caused. (It is for this reason that advocates of increased foreign aid often argue that the rich nations have created the poverty of the poor nations, through forms of economic exploitation that go back to the colonial system.) On the other hand, any consequentialist would insist that we are responsible for all the consequences of our actions; and if a consequence of my spending money on an iPod is that someone dies, I am responsible

for that death. It is true that the person would have died even if I had never existed, but what is the relevance of that? The fact is that I do exist, and the consequentialist will say that our responsibilities derive from the world as it is, not as it might have been.

One way of making sense of the non-consequentialist view of responsibility is to base it on a theory of rights of the kind proposed by John Locke and more recently defended by libertarians like Robert Nozick and Jan Narveson. If everyone has a right to life, and this right is a right *against* others who might threaten my life but not a right *to* assistance from others when my life is in danger, then we can understand the feeling that we are responsible for killing but not for omitting to save. The former violates the rights of others, the latter does not.

Should we accept such a theory of rights? If we build up our theory of rights by imagining, as Locke and Nozick do, individuals living independently from one another in a 'state of nature', it may seem natural to adopt a conception of rights in which as long as each leaves the other alone, no rights are violated. I might, on this view, quite properly have maintained my independent existence if I had wished to do so. So if I do not make you any worse off than you would have been if I had had nothing at all to do with you, how can I have violated your rights? The factual basis of this theory is doubtful. Thomas Pogge challenges it in *World Poverty and Human Rights*, arguing that several features of the world economic order show that we contribute to the impoverishment of some people to our own benefit. To take just one example, we rely on oil and minerals bought from countries ruled by dictators who use the money to enrich themselves or to strengthen their armies and entrench themselves in power. These dictators have no moral right to the wealth that lies beneath the soil of the countries in which they have seized power. The proceeds should go to the people of the country as a whole. The dictators are robbers and murderers, and we are receivers of stolen goods. Our willingness to hand over billions of dollars to dictators in return for oil and mineral rights also creates a huge incentive for anyone who fancies their chances of overthrowing an existing government, and thus increases instability in these countries, which in turn contributes to poverty. (Climate change creates another problem for the view that we are not harming the poor – but that is the topic of the next chapter.)

Even if we put aside such problems with the factual basis of the libertarian argument, we need to ask why we should start from the unhistorical, abstract and ultimately inexplicable idea of a human being living independently. Our ancestors were – like other primates – social beings

long before they were human beings, and they could not have developed the abilities and capacities of human beings if they had not been social beings first. We are not, now, isolated individuals, and we never have been. So why should we assume that rights must be restricted to rights against interference? We might, instead, adopt the view that taking rights to life seriously is incompatible with standing by and watching people die when one could easily save them.

(4) What of the difference in motivation? That a person does not positively wish for the death of another lessens the severity of the blame she deserves, but not by as much as is suggested by our present attitudes to giving aid. The behaviour of the speeding motorist is again comparable, for such motorists usually have no desire at all to kill anyone. They merely want to get somewhere sooner, or they enjoy speeding and are indifferent to the consequences. Despite their lack of malice, those who kill with cars deserve not only blame but also severe punishment.

(5) The difference I have left for last is the most significant. The fact that to avoid killing people is normally not difficult whereas to save all one possibly could save is heroic must make an important difference to our attitude to failure to do what the respective principles demand. Not to kill is a minimum standard of acceptable conduct we can require of everyone. In contrast, to save all one possibly could is not something that can realistically be required, especially not in societies accustomed to giving as little as ours do. Given the generally accepted standards, people who give, say, 10 percent of what they earn to help the poor are more aptly praised for their above-average generosity than blamed for giving less than they might. The appropriateness of praise and blame is, however, a separate issue from the rightness or wrongness of actions. The former evaluates the agent; the latter evaluates the action. Perhaps many people who give 10 percent really ought to give 50 percent, but to blame them for not giving more could be counterproductive. It might make them feel that what is required is too demanding, and if one is going to be blamed anyway, one might as well not give anything at all. That an ethic that puts saving all one possibly can on the same footing as not killing would be an ethic for saints or heroes should not lead us to assume that the alternative must be an ethic that makes it obligatory not to kill but puts us under no obligation to save anyone. There are positions in between these extremes, as we shall see.

Let's summarize the five differences that normally exist between killing and allowing to die in the context of extreme poverty and overseas aid. The lack of an identifiable victim is of no moral significance, though it may play an important role in explaining our attitudes. The idea that

we are directly responsible for those we kill, but not for those we do not help, depends on a questionable notion of responsibility and may need to be based on a dubious theory of rights. Differences in certainty and motivation are ethically significant and show that not aiding the poor is not to be condemned as murdering them; it could, however, be on a par with killing someone as a result of reckless driving, which is serious enough. Finally, the difficulty of completely discharging the duty of saving all one possibly can makes it inappropriate to blame those who fall short of this target in the same way that we blame those who kill; but this does not show that the act itself is less serious. Nor does it excuse those who make no effort to save anyone.

In any case, whereas failing to save a life may not always be ethically on par with deliberate killing, it is clear that how we respond to the existence of both absolute poverty and absolute affluence is one of the great moral issues of our time. So let us consider afresh whether we have an obligation to assist those whose lives are in danger and, if so, how this obligation applies to the present world situation.

THE OBLIGATION TO ASSIST

The Argument for an Obligation to Assist

On my way to give a lecture, I pass a shallow ornamental pond and notice that a small child has fallen in and is in danger of drowning. I look around to see where the parents, or babysitter, are, but to my surprise, I see that there is no one else around. It seems that it is up to me to make sure that the child doesn't drown. Would anyone deny that I ought to wade in and pull the child out? This will mean getting my clothes muddy, ruining my shoes and either cancelling my lecture or delaying it until I can find something dry to change into; but compared with the avoidable death of a child none of these things are significant.

A plausible principle that would support the judgment that I ought to pull the child out is this: if it is in our power to prevent something very bad happening, without thereby sacrificing anything of comparable moral significance, we ought to do it. This principle seems uncontroversial. It will obviously win the assent of consequentialists; but non-consequentialists should accept it too, because the injunction to prevent what is bad applies only when nothing comparably significant is at stake. Thus, the principle cannot lead to the kinds of actions of which non-consequentialists strongly disapprove – serious violations of individual rights, injustice, broken promises and so on. If non-consequentialists

regard any of these as comparable in moral significance to the bad thing that is to be prevented, they will automatically regard the principle as not applying in those cases in which the bad thing can only be prevented by violating rights, doing injustice, breaking promises or whatever else is at stake. Most non-consequentialists hold that we ought to prevent what is bad and promote what is good. Their dispute with consequentialists lies in their insistence that this is not the sole ultimate ethical principle: that it is *an* ethical principle is not denied by any plausible ethical theory.

Nevertheless, the uncontroversial appearance of the principle that we ought to prevent what is bad when we can do so without sacrificing anything of comparable moral significance is deceptive. If it were taken seriously and acted on, our lives and our world would be fundamentally changed. For the principle applies, not just to rare situations in which one can save a child from a pond, but to the everyday situation in which we can assist those living in absolute poverty. In saying this, I assume that absolute poverty, with its hunger and malnutrition, lack of shelter, illiteracy, disease, high infant mortality and low life expectancy, is a bad thing. Additionally, I assume that it is within the power of the affluent to reduce absolute poverty, without sacrificing anything of comparable moral significance. If these two assumptions and the principle we have been discussing are correct, we have an obligation to help those in absolute poverty that is no less strong than our obligation to rescue a drowning child from a pond. Not to help would be wrong, whether or not it is intrinsically equivalent to killing. Helping is not, as conventionally thought, a charitable act that is praiseworthy to do but not wrong to omit. It is something that everyone ought to do.

Set out more formally, this argument would look like this.

First premise:	If we can prevent something bad without sacrificing anything of comparable significance, we ought to do it.
Second premise:	Extreme poverty is bad.
Third premise:	There is some extreme poverty we can prevent without sacrificing anything of comparable moral significance.
Conclusion:	We ought to prevent some extreme poverty.

The first premise is the substantive moral premise on which the argument rests, and I have tried to show that it can be accepted by people who hold a variety of ethical positions.

The second premise is unlikely to be challenged. It would be hard to find a plausible ethical view that did not regard extreme poverty, with the suffering and deaths of both adults and children that it causes, not to mention the lack of education, sense of hopelessness, powerlessness and humiliation that are also its effects, as a bad thing.

The third premise is more controversial, even though it is cautiously framed. It claims only that some extreme poverty can be prevented without the sacrifice of anything of comparable moral significance. It thus avoids the objection that any aid I can give is just 'drops in the ocean', for the point is not whether my personal contribution will make any noticeable impression on world poverty as a whole (of course it won't) but whether it will prevent some poverty. This is all the argument needs to sustain its conclusion, because the second premise says that any extreme poverty is bad and not merely the total amount of extreme poverty. If without sacrificing anything of comparable moral significance we can provide just one family with the means to raise itself out of extreme poverty, the third premise is vindicated.

Nevertheless, some will argue that I can't have any confidence that my donation to an aid organization will save a life or will help people to lift themselves out of extreme poverty. Often these arguments are based on demonstrably false beliefs, such as the idea that aid organizations use most of the money given to them for administrative costs, so that only a small fraction gets through to the people who need it, or that corrupt governments in developing nations will take the money. In fact, the major aid organizations use no more than 20 percent of the funds they raise for administrative purposes, leaving at least 80 percent for the programs that directly help the poor; and they do not donate to governments but work directly with the poor, or with grassroots organizations in developing countries that have a good record of helping the poor.

Measuring the effectiveness of an aid organization by the extent to which it can reduce its administrative costs is, however, a common mistake. Administrative costs include the salaries of experienced people who can ensure that your donation will fund projects that really help the poor in a sustainable, long-term way. An organization that does not employ such people may have lower administrative costs than one that does, but it will still achieve less with your donation.

GiveWell.org is not an aid organization but an organization that seeks hard evidence about which organizations are most effective. It has, for example, compared the cost per life saved of various organizations that work to combat the diseases that kill many of those 8.8 million children

who die each year from poverty-related causes. According to GiveWell, there are several organizations that can save a life for somewhere in the range of $600 to $1200, and on the GiveWell.org Web site, you can see which it ranks most highly. Because you can give to one of the top-ranked organizations, it seems clear that the third premise of the argument is true for people who spend at least a few hundred dollars a year on things they do not really need. They can save a life, or prevent some extreme poverty, without sacrificing anything of comparable moral significance.

I have left the notion of moral significance unexamined in order to show that the argument does not depend on any specific values or ethical principles. On any defensible view of what is morally significant, the third premise will be true for most people living in industrialized nations. Our affluence means that we have income we can dispose of without giving up the basic necessities of life, and we can use this income to reduce extreme poverty. Just how much we will think ourselves obliged to give up will depend on what we consider to be of comparable moral significance to the poverty we could prevent: stylish clothes, expensive dinners, a sophisticated stereo system, exotic holidays, a luxury car, a larger house, private schools for our children and so on. For a utilitarian, none of these is likely to be of comparable significance to the reduction of extreme poverty; and those who are not utilitarians surely must, if they subscribe to the principle of universalizability, accept that at least *some* of these things are of far less moral significance than the extreme poverty that could be prevented by the money they cost. So the third premise seems to be true on any plausible ethical view – although the precise amount of extreme poverty that can be prevented before anything of comparable moral significance is sacrificed will vary according to the ethical view one accepts.

Objections to the Argument

Taking Care of Our Own

Anyone who has worked to increase foreign aid will have come across the argument that we should look after those near us, our families and then the poor in our own country before we think about poverty in distant places.

No doubt we instinctively prefer to help those who are close to us. Few could stand by and watch a child drown; many can ignore the avoidable deaths of children in Africa or India. The question, however, is not what

we usually do, but what we ought to do, and it is difficult to see any sound moral justification for the view that distance, or community membership, makes a crucial difference to our obligations.

Consider, for instance, racial affinities. Should people of European origin help poor Europeans before helping poor Africans? Most of us would reject such a suggestion, and our discussion of the principle of equal consideration of interests in Chapter 2 has shown why we should reject it: people's need for food has nothing to do with their race, and if Africans are in greater need than Europeans, it would be a violation of the principle of equal consideration to give preference to Europeans.

The same point applies to citizenship or nationhood. Every affluent nation has some relatively poor citizens, but absolute poverty is limited largely to the developing nations. In the United States, a family of four is officially classified as poor if they have an annual income of less than $22,000. It can be very difficult to support a family on that income in the United States, but clearly it will take several thousand dollars to make a significant improvement in the lives of people in that situation. In developing countries, on the other hand, it costs less than $1,000 to save the life of a child who would otherwise die from a poverty-related disease, and to double the income of ten families living in extreme poverty would take less than $5,000. (The figure is merely for comparison – I am not suggesting that the best way to reduce poverty is to give money directly to the poor.) Because everyone's resources are limited, it makes sense to use them where they can have the most beneficial impact. Under these circumstances, it would be wrong to decide that only those fortunate enough to be citizens of our own affluent community will share our abundance.

We feel obligations of kinship more strongly than those of citizenship. What kind of parents could give away their last bowl of rice if their own children were starving? To do so would seem unnatural. Indeed, it would be contrary to our nature as biologically evolved mammals with offspring who are dependent on us for many years – but that alone would not show that it was wrong to do so. In any case, we are not faced with that situation but with one in which our own children are well fed, well clothed, well educated and would now like new bikes or more sophisticated computer games. In these circumstances, any special obligations we might have to our children have been fulfilled, and the needs of strangers make a stronger claim on us.

The element of truth in the view that we should first take care of our own lies in the advantage of a recognized system of responsibilities. When families and local communities look after their own poorer members, ties

of affection and personal relationships achieve ends that would otherwise require a large, impersonal bureaucracy. Hence, it would be absurd to propose that from now on we all regard ourselves as equally responsible for the welfare of everyone in the world; but the argument for an obligation to assist does not propose that. It applies only when some are in extreme poverty, and others can help without sacrificing anything of comparable moral significance. To allow one's own kin to sink into extreme poverty would be to sacrifice something of comparable significance; and well before that point had been reached, the breakdown of the system of family and community responsibility would be a factor to weigh the balance in favour of a modest preference for family and community. This modest degree of preference is, however, decisively outweighed by existing discrepancies in wealth and property.

Property Rights

Do people have a right to private property, a right that contradicts the view that they are under an obligation to give some of their wealth away to those in absolute poverty? According to some theories of rights, people who have acquired their property without the use of unjust means like force and fraud may be entitled to great wealth and every conceivable luxury while others starve. This individualistic conception of rights is in contrast to other views, like the Christian doctrine that holds that property exists for the satisfaction of human needs; and so, as Thomas Aquinas wrote, 'whatever a man has in superabundance is owed, of natural right, to the poor for their sustenance'. A socialist would also, of course, see wealth as belonging to the community rather than the individual; whereas utilitarians, whether socialist or not, would be prepared to override property rights to prevent great evils.

Does the argument for an obligation to assist others therefore presuppose one of these other theories of property rights and reject the idea of a strong individual right to property? Not necessarily. A theory of property rights can insist on our *right* to retain wealth without pronouncing on whether the rich *ought* to give to the poor. Robert Nozick, for example, rejected the use of compulsory means like taxation to redistribute income, but he suggested that we can achieve the ends we deem morally desirable by voluntary means. So Nozick would have rejected the claim that rich people have an 'obligation' to give to the poor, insofar as this implies that the poor have a right to our aid, but could have agreed that giving is something we ought to do and failing to give – though

within one's rights – is wrong, for there is more to an ethical life than respecting the rights of others.

The argument for an obligation to assist can survive, with only minor modifications, even if we accept an individualistic theory of property rights. In any case, however, I do not think we should accept such a theory. It leaves too much to chance to be an acceptable ethical view. For instance, many of those whose forefathers happened to inhabit some sandy wastes around the Persian Gulf are now fabulously wealthy, because oil lay under those sands; whereas many of those whose forefathers settled on better land south of the Sahara live in extreme poverty, because of drought and bad harvests. Can this distribution be acceptable from an impartial point of view? If we imagine ourselves about to begin life as a citizen of either Kuwait or Chad – but we do not know which – would we accept the principle that citizens of Kuwait are under no obligation to assist people living in Chad?

Population and the Ethics of Triage

Leaving it to the Government

We often hear that foreign aid should be a government responsibility and not left to private charity. Giving privately, it is said, allows the government to escape its responsibilities. If we give, the government won't see the need to do so.

Because increasing government aid is the surest way of significantly increasing the total amount of aid given, I would agree that the

governments of affluent nations should give more aid than they give now – as long as it is given for projects that effectively help those in extreme poverty. Less than 25 cents in every $100 dollars of gross national income is a scandalously small amount for a nation as wealthy as the United States to give to relieve extreme poverty in the world's poorest countries – and that figure includes both government aid and non-government charitable donations. Even the official UN target of 0.7 percent seems much less than affluent nations can and should give – though it is a target few have reached. But is this a reason against each of us giving as much as we can through voluntary agencies? To believe that it is seems to assume that the more people there are who give through voluntary agencies, the less likely it is that the government will do its part. Is this plausible? The opposite view – that if no one gives voluntarily, the government will assume that its citizens are not in favour of overseas aid and will cut its programme accordingly – is more reasonable. In any case, unless there is a definite probability that by refusing to give we would be helping to bring about an increase in government assistance, refusing to give privately is wrong for the same reason that shutting off aid because of the risks of overpopulation is wrong: it is a refusal to prevent a definite evil for the sake of a very uncertain gain. The onus of showing how a refusal to give privately will make the government give more is on those who refuse to give.

This is not to say that giving privately is enough. As active concerned citizens, we should campaign for entirely new standards for both public and private aid. We should also work for fairer trading arrangements between rich and poor countries, including an end to rich nations paying subsidies to their agricultural producers that make it impossible for poor countries to compete in global markets. Perhaps it is more important to be politically active in the interests of the poor than to give to them oneself – but why not do both? Unfortunately, many use the view that aid is the government's responsibility as a reason against giving but not as a reason for being politically active.

Too High a Standard?

The final objection to the argument I have given for an obligation to assist is that it is too demanding; it sets a standard so high that only a saint could attain it. This objection comes in at least three versions. The first maintains that, human nature being what it is, we cannot achieve so high a standard; and because it is absurd to say that we ought to do what we cannot do, we must reject the claim that we ought to give so much. The second version asserts that even if we could achieve so high a standard, to do so would be undesirable. The third version of the objection is that to set so high a standard is undesirable because it will be perceived as too difficult to reach and will discourage many from even attempting to do so.

Those who put forward the first version of the objection often make observations about human nature. They point out that we all are much more concerned about our own interests, and those of our immediate family, than we are about the interests of strangers. That is, they may add, because we have evolved from a natural process in which those with a high degree of concern for their own interests, or the interests of their offspring and kin, tended to leave more descendants in future generations than those who were not so concerned with their own interests or those of their kin. Thus, the biologist Garrett Hardin has argued, in support of his 'lifeboat ethics', that altruism can only exist 'on a small scale, over the short term, and within small, intimate groups'; and Richard Dawkins has written, in his provocative book *The Selfish Gene*: 'Much as we might wish to believe otherwise, universal love and the welfare of the species as a whole are concepts which simply do not make evolutionary sense.'

I have already noted, in discussing the objection that we should first take care of our own, the very strong tendency for partiality in human beings. Our preference for our own interests, and those of our close kin, over the interests of strangers is no doubt a natural outcome of the evolutionary process. What this means is that we would be foolish to expect widespread conformity to a standard that demands impartial concern, and for that reason it would scarcely be appropriate or feasible to condemn all those who fail to reach such a standard. Yet to act impartially, though it might be very difficult, is not impossible. The commonly quoted maxim '*ought* implies *can*' does not apply here. That maxim is a reason for rejecting such moral judgments as, 'You ought to have saved all the people from the sinking ship', when in fact if you had taken one more person into the lifeboat, it would have sunk and you would not have saved any. In that situation, it is absurd to say that you ought to have done what you could not possibly do. When we have money to spend on luxuries and others are starving, however, it is clear that we *can* all give much more than we do give, and we can therefore all come closer to the impartial standard proposed in this chapter. Nor is there, as we approach closer to this standard, any barrier beyond which we cannot go.

A remarkable illustration of what is possible for a family to do began in Atlanta, Georgia, in 2006 when the car in which Kevin Salwen was driving his fourteen-year-old daughter Hannah was halted by a stoplight. On one

side Hannah saw a gleaming Mercedes coupe, and on the other she saw a homeless man. 'You know, Dad,' she said, pointing, 'if *that* man had a less nice car, that man *there* could have a meal.' That started a conversation that continued at home. Instead of scoffing at the idea, Hannah's mother challenged her: 'What do you want to do: sell our house, move into one half the size, give up your room?' Over a series of family discussions, the Salwens, a well-off family of four, decided to do just that: sell their home, give half the money they received for it to the poor and, with the other half, buy a smaller home. Friends thought they were crazy, but they were confident that they were doing the right thing. As a result, they were able to give more than $800,000 to help rural villagers in Ghana lift themselves out of poverty. Many people would consider moving to a smaller home a sacrifice, but Kevin Salwen says that even in terms of self-interest, it made sense: 'Giving away half of something we had too much of (our house) brought us a togetherness, trust and joy we never had.'

Admittedly, the Salwens' decision still left them comfortably off. They could have given more without sacrificing anything comparable in significance to the lives that they could, by giving even more, have saved. So this example does not demonstrate that this standard is achievable, but it does show how a family can break through barriers that most of us take for granted. Zell Kravinsky pushed those barriers even further. After making more than $40 million though canny real estate investments, he gave away almost all of it, living with his family in a modest suburban home. Then, on learning that people die from kidney disease while waiting for a transplant to become available, and studying research showing that the chances of anyone needing both kidneys are as low as 1 in 4000, he went to a city hospital that served mostly African Americans and donated one of his kidneys to a stranger. With examples like these, we cannot say that the impartial standard is mistaken because it is impossible for us – for anyone one of us, individually – to achieve it. We don't really know how far in the direction of impartiality it is possible to go. Unlike the Salwens or Zell Kravinsky, most people never try.

The second version of the objection has been put by several philosophers during the past decade, among them Susan Wolf in a forceful article entitled "Moral Saints". Wolf argues that if we all took the kind of moral stance defended in this chapter, we would have to do without a great deal that makes life interesting: opera, gourmet cooking, elegant clothes and professional sport, for a start. The kind of life we come to see as ethically required of us would be a single-minded pursuit of the

overall good, lacking that broad diversity of interests and activities that, on a less demanding view, can be part of our ideal of a good life for a human being. To this, however, one can respond that although the rich and varied life that Wolf upholds as an ideal may be the most desirable form of life for a human being in a world of plenty, it is wrong to assume that it remains a good life in a world in which buying luxuries for oneself means accepting the continued avoidable suffering of others. A doctor faced with hundreds of injured victims of a train crash can scarcely think it defensible to treat fifty of them and then go to the opera, on the grounds that going to the opera is part of a well-rounded human life. The life-or-death needs of others must take priority. Looking at the world as a whole, and our ability to make a difference, we are like the doctor in that we live in a time when we all have an opportunity to help to mitigate a disaster.

Associated with this second version of the objection is the claim that an impartial ethic of the kind advocated here makes it impossible to have serious personal relationships based on love and friendship. These relationships are, of their nature, partial. We put the interests of our loved ones, our family and our friends ahead of those of strangers. If we did not do so, would these relationships survive? I have already indicated, in the response I gave when considering the objection that we should first take care of our own, that there is a place within an impartially grounded moral framework for recognising some degree of partiality for kin, and the same can be said for other close personal relationships. Clearly, for most people personal relationships are among the necessities of a flourishing life, and to give them up would be to sacrifice something of great moral significance. Moreover, for most people, to give up such relationships would diminish, not only their happiness and their mental health, but also their effectiveness as an agent of change. Hence, no such sacrifice is required by the principle for which I am here arguing.

The third version of the objection asks: might it not be counterproductive to demand that people give up so much? Might not people say, 'As I can't do what is morally required anyway, I won't bother to give at all'? If, however, we were to set a more realistic standard, people might make a genuine effort to reach it. Thus, setting a lower standard might actually result in more aid being given.

It is important to get the status of this third version of the objection clear. Its accuracy as a prediction of human behaviour is quite compatible with the argument that we are obliged to give to the point at which by giving more we sacrifice something of comparable moral significance to

what we achieve by our donation. What would follow from the objection is that public advocacy of this standard of giving is undesirable. It would mean that, in order to do the maximum to reduce extreme poverty, we should advocate a standard lower than the amount we think people really ought to give. Of course we ourselves – those of us who accept the original argument, with its higher standard – would know that we ought to do more than we publicly propose people ought to do, and we might actually give more than we urge others to give. There is no inconsistency here, because in both our private and our public behaviour we are trying to do what will most reduce extreme poverty.

For a consequentialist, this apparent conflict between public and private morality is always a possibility and not in itself an indication that the underlying principle is wrong. The consequences of a principle are one thing, the consequences of publicly advocating it another. A variant of this idea is already acknowledged by the distinction between the intuitive and critical levels of morality, of which I have made use in previous chapters. If we think of principles that are suitable for the intuitive level of morality as those that should be generally advocated, these are the principles that, when advocated, will give rise to the best consequences. Where aid is concerned, they will be the principles that lead to the largest amount being given by the affluent to the poor – as long as the money is given to an organization that will use it with maximum effectiveness.

Is it true that the standard set by our argument is so high as to be counterproductive? There is not much evidence to go by, but discussions of the argument with students and others have led me to think it might be. On the other hand, the conventionally accepted standard – a few coins in a collection tin when one is waved under your nose – is obviously far too low. What level should we advocate? In my book *The Life You Can Save* – and on the corresponding website, www.thelifeyoucansave.com – I have suggested a progressive scale, like a tax scale. It begins at just 1 percent of income; and for 90 percent of taxpayers, it does not require giving more than 5 percent. This is therefore an entirely realistic amount, and one that people could easily give with no sacrifice – and indeed, often with a personal gain, because there are many psychological studies showing that those who give are, as the Salwen family found, happier than those who do not. I do not really know if the scale I propose is the one that will, if widely advocated, achieve the greatest total amount donated, but I calculated that if everyone in the affluent world gave according to that scale, it would raise $1.5 trillion each year – which is eight times what the United Nations task force headed by the economist Jeffrey

Sachs calculated would be needed to meet the Millennium Development Goals set by the leaders of all the world's nations when they met at the UN Millennium Development Summit in 2000. Those goals included reducing by half the proportion of the world's people living in extreme poverty and the proportion of people who suffer from hunger, as well as reducing by two-thirds the death toll among children under five years old – thus saving six million lives every year – and enabling children everywhere to have a full course of primary schooling.

This surprising outcome – that if everyone with abundance were to contribute to the effort to reduce extreme poverty and all that goes with it, the amount each of us would need to give would be quite modest – shows that the argument with which this chapter began is demanding only because so few of those with the ability to help the poor are doing anything significant to help them. We do not need to transfer half or a quarter or even a tenth of the wealth of the rich to the poor. If few are helping, those few have to cut very deep before they get to the point at which giving more would involve sacrificing something of comparable moral significance to the life saved by their gift. If we all, or even most of us, gave according to the much more modest scale I have suggested, none of us would have to give up much. That is why this is a suitable standard for public advocacy. What we need to do is to change our public ethics so that for anyone who can afford to buy luxuries – and even a bottle of water is a luxury if there is safe drinking water available free – giving something significant to those in extreme poverty becomes an elementary part of what it is to live an ethical life.