

Distributing Responsibilities*

DAVID MILLER

Politics, Nuffield College, Oxford

IN this article I examine a question that arises frequently in moral and political debate, but has not to my knowledge been examined in much depth by philosophers. Our world contains all too many instances of deprived or suffering people—people whose basic rights to security, or subsistence, or health care are not being protected, and who as a result are in no position to live minimally decent lives. Nearly all of us believe that this is a situation that demands a remedy: someone should provide the resources to end the suffering and deprivation.¹ The problem does not lie here, but in deciding which particular agent or agents should put the bad situation right. Very often there are many agents who could act in this way. The issue is how to identify one particular agent, or group of agents, as having a particular responsibility to remedy the situation. For unless we can do this, there is a danger that the suffering or deprivation will continue unabated, even though everyone agrees that it is morally intolerable, because no-one is willing to accept the responsibility to step in and relieve it.

For an example of the problem I have in mind, consider the current plight of Iraqi children who are malnourished and lack access to proper medical care. No one doubts that their condition is a very bad one, nor is it difficult to grasp what would be needed to remedy it. But who has the responsibility for putting it right? Is it the United Nations, and more especially the Western powers, who on the one hand have contributed to the situation, or so it is alleged, by imposing economic sanctions on Iraq, and on the other are well placed to supply the necessary food and medical aid? Or is it Saddam Hussein and his henchmen, who have diverted a large percentage of Iraq's GNP to military expenditure, and who have deliberately (it is alleged) prevented aid from reaching poor families in an attempt to have the sanctions lifted? Or does responsibility lie with the Iraqi people as a whole, on the grounds that each nation has a duty to look after its own, which in this case might involve taking direct action to overthrow the current brutal regime?

*Earlier versions of this article were presented to the Nuffield Political Theory Workshop and to the All Souls Seminar in Political Philosophy. I should like to thank both audiences for their helpful comments. I should particularly like to thank Jerry Cohen, Richard Dagger, Cécile Fabre and Bob Goodin, plus three anonymous referees for this *Journal*, for their extensive written comments on earlier drafts of the article.

¹I say 'nearly all' because there may be philosophers with libertarian instincts who hold that deprivation and suffering call for remedy only when they are the result of some agent violating the victims' rights. I examine this position briefly later in the article.

I shall call this the problem of remedial responsibility. To be remedially responsible for a bad situation means to have a special obligation to put the bad situation right, in other words to be picked out, either individually or along with others, as having a responsibility towards the deprived or suffering party that is not shared equally among all agents. The problem is to find a principle, or set of principles, for assigning such responsibilities which carries moral weight, so that we can say that agents who fail to discharge their remedial responsibilities act wrongly and may properly be sanctioned. (What form the sanctioning may take will vary from case to case, and it is not part of my brief here to pursue this question. I mention sanctions simply to underline the point that when we are arguing about where the responsibility for remedying a bad situation should fall, we mean our answer to have some teeth.) In other words, the problem is: what connects a particular agent A to a particular patient P in such a way that A is singled out as having a remedial responsibility towards P that others, in general, do not have? Note that the agents in question may be individual people, but, as the example I gave above illustrates, they may also be collectives of various kinds—governments, states, corporations, even those amorphous entities called nations. Assigning responsibility to these collective bodies raises additional questions that I cannot address here.² Instead I shall assume that such assignments of collective responsibility are both meaningful and justifiable, and focus on the issue of how they should be distributed. In exploring this issue, I shall often refer to individual agents and patients, because we are likely to have better formed judgements in these cases, but eventually our analysis should be applicable to collective cases too.

Because the problem of distributing responsibilities is so urgent, human societies have evolved mechanisms whereby they are formally assigned to individual people or to institutions. If we ask who is responsible for safeguarding this particular battered child, the answer is likely to be the social worker who has been assigned to the case. I am not concerned here with such formally assigned responsibilities, but with the underlying principles that should guide us when we are in a position to make formal assignments, and that we should appeal to directly when no formal assignments have been made. Very often, in fact, the problem arises precisely because of the lack of any institutional mechanism that can assign responsibilities formally—the international arena is replete with examples. We may believe that we should move towards a situation where for any group of deprived or suffering people there is some agency that has formally been assigned the responsibility to remedy their condition. But clearly that day is a very long way off, and meanwhile the best we can do is to lay out some principles for distributing responsibilities that we hope will command widespread agreement. That is the task of the present article.

²I have tackled some of them in a companion paper, 'Holding nations responsible'.

I. PRINCIPLES OF REMEDIAL RESPONSIBILITY

Perhaps the most obvious solution to our problem is to say that agents should be held remedially responsible for situations when, and to the extent that, they were responsible for bringing those situations about. In other words we look to the past to see how the deprivation and suffering that concerns us arose, and having established that, we are then able to assign remedial responsibility. In the case of the Iraqi children, for instance, we need to know why they are malnourished and sick: who is responsible for bringing about this state of affairs?

Clearly this answer invokes a different sense of ‘responsibility’ from the one that directly concerns us. But which sense? Unfortunately few concepts in moral and political philosophy are more slippery than that of responsibility, and it is a fair bet that real debates on issues such as the plight of Iraqi children become muddled as the protagonists slip from one meaning of responsibility to the next.³ So we need to draw some distinctions, and in particular a distinction between *causal* and *moral* responsibility.

To say that an agent is causally responsible for some state of affairs is simply to highlight the causal role played by the agent in the genesis of that state of affairs. Here I rely on what Hart and Honoré have called the common sense understanding of causation according to which when we say that C caused E we are singling out C as one among a potentially large number of antecedent conditions for E’s occurrence, distinguished from the other conditions by virtue of its abnormality (and also, in many cases, by virtue of its being a deliberate human action).⁴ For an example of ‘bare’ causal responsibility—causal responsibility that is not accompanied by moral responsibility—consider the case where I am walking along the pavement, taking ordinary care, but trip over a raised paving stone, knock down the person in front of me and injure him. Then I am causally responsible for the injury, but not morally responsible, because I have done

³It is worth citing a story invented by Hart to illustrate the slipperiness of the concept: ‘As captain of the ship, X was responsible for the safety of his passengers and crew. But on his last voyage he got drunk every night and was responsible for the loss of the ship with all [others] aboard. It was rumoured that he was insane, but the doctors considered that he was responsible for his actions. Throughout the voyage he behaved quite irresponsibly, and various incidents in his career showed that he was not a responsible person. He always maintained that the exceptional winter storms were responsible for the loss of the ship, but in the legal proceedings brought against him he was found criminally responsible for his negligent conduct, and in separate civil proceedings he was held legally responsible for the loss of life and property. He is still alive and he is morally responsible for the deaths of many women and children’. (H. L. A. Hart, *Punishment and Responsibility: Essays in the Philosophy of Law* (Oxford: Clarendon Press, 1968), p. 211.)

⁴See H. L. A. Hart and T. Honoré, *Causation in the Law*, 2nd edn (Oxford, Clarendon Press, 1985), ch. 2 for this account. Their analysis raises many questions, not least the problem that the distinction between causes and (mere) conditions depends upon the perspective from which we are looking at any given event, which in turn determines what will be regarded as normal background conditions and what will be regarded as an abnormal intervening circumstance. This in turn raises the possibility that causal judgements are in part influenced by moral considerations that determine the perspective from which our causal enquiry is launched. I return briefly to this point later in the article. Here, I assume that we have a workable notion of causal responsibility that is distinct from moral responsibility, as illustrated by the example given in the text, even if there are ineliminable practical concerns that lie behind our judgements of causal responsibility.

nothing that attracts moral praise or blame. My tripping is simply the unusual feature in the case that accounts for the unfortunate injury to the pedestrian.

Moral responsibility, on the other hand, involves an appraisal of the agent's conduct. In the cases that particularly concern us, the agent's role in bringing about the outcome must be such that it leaves the agent liable to moral blame. That in turn requires us to ask questions such as whether the agent intended the outcome, whether he foresaw it, whether his behaviour violated some standard of reasonable care, and so forth. As the example above shows, moral responsibility is in one respect a narrower notion than causal responsibility, since there will be many cases in which someone's conduct is perfectly innocent but it just so happens that something he does is the main causal factor in injuring another. But in another respect it may be wider, for instance if I negligently fail to take steps to prevent something from occurring. Suppose I take my son Jamie and his friend Nick to the park to play, and in the course of some rather boisterous game Jamie manages to break Nick's arm. Meanwhile I am sitting on the bench with my head buried in a newspaper and fail to notice what is going on. Then I may be morally responsible for Nick's arm getting broken, even though it is clearly Jamie who is causally responsible according to the criteria suggested above (whether Jamie is also morally responsible will depend on the details of the case—essentially whether it is reasonable to expect a boy of that age to foresee the likely outcome of the rough-and-tumble that is taking place). I am morally responsible because I have failed in my duty to take care of the boys, something that I assumed when I offered to take them down to the park. I can properly be blamed for not preventing the broken arm.

Having clarified the distinction between causal and moral responsibility, we can now ask whether either of these yields an adequate principle for assigning remedial responsibility. The appeal of causal responsibility is straightforward. If A is the cause of P's deprived condition—and this appears to be something that can be established empirically—then what is more obvious than to hold A responsible for remedying it? If he was the sole cause, then remedial responsibility is his alone; if several agents together caused P's condition, then remedial responsibility should be distributed in proportion to causal responsibility. But on closer inspection the causal principle faces a number of damaging objections.

The first is simply that there are many cases in which no identifiable A has caused P's condition, and yet we would be reluctant to say that no-one has any remedial responsibility towards P. Prominent here are instances in which P's condition results from natural causes—P starves because of crop failure or is stricken with cancer. Of course it is always possible in such instances to specify forms of action which, had certain agents taken them, would have prevented the condition from emerging. The crop failure would not have occurred if company A had installed an irrigation system; the cancer would not have occurred if health authority B had decided to invest an extra £X million in preventative medicine.

But there are an infinite number of such counterfactuals, and so they will not, in general, identify any particular agents as causally responsible for P's condition. When a particular counterfactual seems relevant, that is because the agent it describes has already been identified as bearing responsibility for P's condition. Thus if company A had contracted to install an irrigation system in P's neighbourhood, it *then* becomes appropriate to single out A's inaction as the cause of the crop failure. But in the absence of any such agreement, or other special link between A and P, there is no reason to distinguish A's failure to install an irrigation system from B's failure to supply fertiliser, from C's failure to make available genetically modified seeds, and so on—the list containing all those actions the performance of any one of which would have prevented the crop failure.⁵

A second, related, difficulty arises when there are several agents whose actions can be plausibly linked to P's condition. Here there seems to be no merely empirical way of dividing up causal responsibility as a basis for assigning remedial responsibility.⁶ Return to the case of the deprived Iraqi children. Suppose it is true, as seems plausible, that if the UN had not decided to impose economic sanctions, more money would have been available to fund health and social security in Iraq. Suppose it is also true that if Saddam had decided to cut military expenditure, enough would have been left, even with the sanctions, to prevent the destitution. Both the UN and Saddam can then be described as causally responsible for the sufferings of the children. But how should we apportion responsibility between them? In answering this question, we cannot apparently avoid making moral appraisals of the relevant agents' conduct, especially examining how far their behaviour was justified. If we think that Saddam's regime posed a serious threat to neighbouring countries, and therefore that UN sanctions were justified, then we will single out Saddam's policy decisions as the cause of the children's suffering. If, by contrast, we see Iraq as a vulnerable regime surrounded by enemies, and therefore as justified in arming itself in self-defence, we shall lay causal responsibility at the door of the UN. Causal attributions are being determined by normative assumptions about justified behaviour. But in that case we no longer have a causal principle in our original sense—the ethical question about who bears remedial responsibility for P's condition is no longer being answered just by looking empirically at who brought that condition about.

⁵Why not then treat *all* of these as in equal measure responsible for P's condition? If we do this then we lose the distinctive purpose of responsibility assignments as I understand them, which is to identify one or more agents who are under some special obligation to relieve P's condition, and who therefore can properly be put under pressure to act. In other words, responsibility loses its practical force if it is diffused among all those agents of whom it is true that they *might* have acted in such a way that P's condition did not occur. I return later to the reasons we have for wanting specific assignments of responsibility, not diffuse ones.

⁶This can apply even where A and P herself are the only agents involved, as pointed out in S. Perry, 'The moral foundations of tort law', *Iowa Law Review*, 70 (1992), 449–514, at pp. 463–4.

Finally, questions about justification appear to arise even in cases where there is only one agent who can plausibly be described as the cause of P's condition. For sometimes A may act in a way that is harmful to P, and yet bear no remedial responsibility for the harm he has caused, because we judge A's behaviour to be legitimate. Suppose for instance that in a market setting A drives P out of business by offering a better service to customers, then provided he deals fairly we do not think that he bears any remedial responsibilities towards P. Or A may cast P into depression by marrying the love of his life. So it seems that the causal principle taken by itself cannot explain our remedial responsibilities. It falls down where no particular A can be identified as the cause of P's condition, and it also falls down where A is certainly causally responsible, in whole or in part, for the harm done to P, but A's behaviour appears justifiable, and therefore does not bring remedial responsibilities in its train.

Yet it is interesting to notice that even innocent causation may place *some* special responsibility upon the agent. Return to the case where by simple mischance I stumble and knock down a pedestrian. Everyone in the vicinity is under some obligation to help him to his feet and make sure that he is not badly hurt; yet we believe that the responsibility is in the first place mine, so I have the primary obligation to act.⁷ I did nothing wrong—indeed I could not help doing what I did—yet having done it, having been the cause of P falling to the ground, I seem to be linked to him more strongly than B who just happened to be passing by. (In the same way, it seems appropriate that I should apologise, or at least express regret for what has happened, even though I am not at all to blame for the event.⁸) This may seem mysterious; indeed some may find it disturbing that we can apparently incur responsibilities just by doing something as unintentional and innocent as walking along the street. I do not want to suggest that bare causation of this kind plays more than a minor part in distributing responsibilities. But the fact that it plays any part at all may help us in searching for the correct theory.

So let us now consider the alternative principle that A should be held remedially responsible for P's condition insofar as he is *morally responsible* for its occurrence, in the sense that carries with it ascriptions of fault and blame. This principle seems to capture nicely what is at stake in our original example: when we ask who is responsible for the plight of Iraqi children, we appear to be asking who is morally responsible for bringing about their condition of malnutrition and ill health, whether by acting wrongly (the UN and Saddam) or by failing to act as duty requires (the Iraqi people as a whole). We have seen already that moral

⁷Someone might argue here that we assign responsibility in this way because the person who knocked the pedestrian over is also likely to be the person best placed (by virtue of proximity) to take care of him afterwards. However if we think of the accident occurring in a crowded street, it seems that any advantage of this kind will be negligible, whereas our sense that it is the person who caused the fall who bears the primary responsibility is quite strong.

⁸Cf. T. Honoré, 'Responsibility and luck', *Law Quarterly Review*, 104 (1988), 530–53, at pp. 544–5.

responsibility is narrower than causal responsibility in some ways and wider in others, and it looks as though in both respects this enables it to fit better with our considered judgements about remedial responsibility. There are cases in which people are causally but not morally responsible for the outcome of their actions—for instance those in which the chain of events connecting action and outcome is long and tortuous, so that the agent could not have reasonably foreseen the final result of her action—and here we are unlikely to hold them remedially responsible for the harm that may ensue. Equally, in cases where we judge people to be morally responsible for the occurrence of a harm, even though their causal role in bringing it about was merely a negative one—cases like the delinquent parent who is reading a newspaper when he should have been watching out for children—we *do* hold them remedially responsible. So at first glance it seems that the moral responsibility principle is going to perform better than the causal principle in explaining our remedial responsibilities.

But that is not to say that it explains everything. To begin with, it cannot explain why causation alone sometimes seems sufficient to generate remedial responsibilities—as in the example of the innocent pedestrian. Nor can it deal happily with cases of justified, but harmful, behaviour. For instance, suppose that A must quickly find a certain drug to save Q's life, and the only way he can do this is to steal some from P, who needs the drug too, but less urgently. We think that he should steal the drug from P, but that he then has a remedial responsibility to P to replace what he has taken.⁹ But he is not morally responsible for harm suffered by P in the intervening period in the sense that leaves him liable to attributions of blame; provided he goes on to discharge his remedial responsibilities, his conduct is not faulty.

There is, however, an ambiguity in the meaning of moral responsibility that needs to be addressed at this point. As defined above, it is linked conceptually to liability to blame: to say of A that he is morally responsible for state S is to say that he has contributed to the bringing about of S in such a way as to incur blame. But one may sever this link and use a broader concept of moral responsibility according to which people are to be held morally responsible for the results of their actions, so long as these actions themselves satisfy certain conditions of intentionality, voluntariness etc., without implying that they are blameworthy for what they did.¹⁰ On this second view, A is morally responsible for

⁹A parallel case, involving a backpacker caught in a blizzard who breaks into an unoccupied hut and uses up what he finds there to keep himself alive, was introduced by Joel Feinberg in 'Voluntary euthanasia and the inalienable right to life', *Philosophy and Public Affairs*, 7 (1978), 93–102, and has subsequently been widely discussed—see, for instance, J. J. Thomson, 'Rights and compensation', *Rights, Restitution and Risk* (Cambridge, Mass.: Harvard University Press, 1986) and L. Lomasky, 'Compensation and the bounds of rights', *Nomos XXXIII: Compensatory Justice*, ed. J. W. Chapman (New York: New York University Press, 1991).

¹⁰It is not necessary here to spell this broader, non-appraisive concept of moral responsibility out in detail, except to say that it still remains distinct from causal responsibility: on the broader concept A is morally responsible for the outcome of all the actions he deliberately performs, but not, for instance, for accidentally knocking down a pedestrian as a result of stumbling himself.

the effects on P of stealing the drug, because he took the drug deliberately and freely, even though he was fully justified in acting as he did. The broader view, then, preserves the connection between moral responsibility and remedial responsibility in the drug case. But it runs into difficulties in cases where justified, but harmful, conduct appears not to bring remedial responsibilities with it—cases like that in which A engages in fair competition with P in a market setting. Suppose that A sets up a shop close to P's and by legitimate means attracts most of P's customers, then A is morally responsible for P's loss of earnings, on the broader view. Even if A did not intend the precise outcome that occurred, he went into business deliberately, and could reasonably have foreseen as one result of his competitive behaviour that P would be damaged. Yet we do not feel that he owes P anything by way of remedy so long as he has acted fairly.

So neither of the two concepts of moral responsibility we have considered—neither the narrow concept linked to blame, nor the wide, morally neutral, concept—gives us the link we are looking for between moral and remedial responsibilities. The narrow concept fails to explain why there are remedial responsibilities in the drug case; but the wider concept, which holds people responsible for the results of all their voluntary actions, whether blameworthy or not, implies that there are remedial responsibilities in the shopkeeper case too.¹¹ Neither concept appears to hit the target precisely.

The biggest problem with the moral responsibility principle, however—one that it shares with the causal principle—is that it looks too exclusively to the past in assigning remedial responsibilities. The question it asks is always 'Who is responsible for bringing this bad situation about?' and never, for instance, 'Who is best placed to put it right?' One obvious defect of the principle, therefore, is that it has nothing to say when the morally responsible agent proves to be incapable of discharging her remedial responsibilities—for instance when she is dead or incapacitated. Unless we want to say that remedial responsibilities disappear when we cannot find an agent who is both morally responsible for the situation in question and capable of remedying it, the principle remains incomplete.

Taking our cue from the last paragraph, consider next the principle of *capacity*, which holds that remedial responsibilities ought to be assigned according to the capacity of each agent to discharge them. The rationale for this is obvious enough. If we want bad situations put right, we should give the

¹¹Can the moral responsibility principle be saved by adding in further conditions that differentiate the two cases? One difference between them is that P has a right to the drug that is stolen in the first case, but no right to his customers' patronage in the second case. This suggests the following: A is remedially responsible for P's condition if and only if he is morally responsible for a rights-violation that led to that condition. However on closer inspection this turns out to be too restrictive. Suppose, for instance, that A publishes an unfair review of P's book, damaging his career. We may think that he has a responsibility to offset the damage (supposing there is some way he can do this) without believing that P has a right that has been infringed by the publication of the review. The wrongness of A's action together with the harm suffered by P seem sufficient in this case to generate a remedial responsibility without the invocation of a rights-violation.

responsibility to those who are best placed to do the remedying. If there is a bather in trouble off the beach, and it makes most sense for one person to undertake the rescue, then we should choose the strongest swimmer. In other circumstances we may want to share responsibility among A, B, C . . . in proportion to their respective capacities to rectify P's condition, as suggested by the slogan 'From each according to his abilities, to each according to his needs'.

On closer inspection, however, the capacity principle seems to blend together two different factors which may not always point in the same direction. One has to do with the *effectiveness* of different agents in remedying the situation; the other has to do with the *costs* they must bear in the course of doing so. The strongest swimmer may also be fearful (so that although he is an effective rescuer, the rescue causes him considerable distress)—or perhaps he simply dislikes the kind of attention that goes along with a successful rescue. If A is slightly stronger than B, but A's costs are also much higher, is it obviously the right solution to hold A responsible for rescuing P? On the moral responsibility view costs do not present ethical problems: if A has harmed P through some action of his, then he should remedy the harm regardless of the cost to himself; requiring that he should pay the cost here is justifiable in the light of what he has done.¹² But to apply the capacity principle, it seems, we have to begin by weighing effectiveness against cost to determine whose capacity is the greatest in the morally relevant sense.¹³

Another problem with the capacity principle is that, by focusing attention entirely on agents' present capacity to remedy some harm, it neglects to ask how variations in capacity have arisen. And so it is vulnerable to the grasshopper and ant objection: assuming that the grasshoppers *could* have spent the summer gathering food for the winter rather than singing, we may wonder whether the ants, who now have the capacity to help the famished grasshoppers by virtue of their earlier diligence, have a remedial responsibility to do so. Perhaps the destitution of the grasshoppers imposes some residual responsibility on the ants, but not as much as if, say, the grasshoppers' store of food had been washed away by unexpected rain. The capacity view cannot explain why remedial responsibility is stronger in one case than in the other. This is the weakness that accompanies its strength: its exclusive focus on the present necessarily blinds it to historical considerations.

Finally, we need to ask whether capacity alone—the simple physical ability to remedy P's bad condition—is sufficient to generate remedial responsibility in the absence of a stronger link between A and P. Return to the case of the drowning

¹²This claim requires some fine-tuning. Even if A is unquestionably at fault in acting as he did, he may be liable only for the effects of his action that a reasonable person would have foreseen, not for consequences that arise in peculiar ways or through the intervention of other actors. For discussion of this principle in the context of tort law, see Hart and Honoré, *Causation in the Law*, ch. 9, and A. Ripstein, *Equality, Responsibility and the Law* (Cambridge: Cambridge University Press, 1999), ch. 4.

¹³We could of course simply define capacity as effectiveness, in which case we would have to allow in, as a separate and competing principle, the principle that remedial responsibilities should be assigned to the agents who would bear the least costs in discharging them.

bather, and notice that we single out the strongest swimmer from among those already gathered on the beach. Perhaps, then, we appeal to capacity only after identifying a set of agents whose relationship to P is such that they already bear a special responsibility towards him; capacity is used to pick out one particular agent from the set.¹⁴

That thought suggests a fourth principle for distributing responsibilities which we might call the communitarian principle, using 'community' here in a fairly loose sense to capture special ties of various kinds such as those that exist within families, collegial groups of various kinds, nations, and so forth. The claim is that when people are linked together by such ties, whether arising from shared activities and commitments, common identities, common histories, or other such sources, they also (justifiably) see themselves as having special responsibilities to one another, responsibilities that are greater than those they have towards humanity at large; and this in particular imposes special responsibilities towards any member of the relevant community who is harmed or in need.

The great merit of the communitarian principle is that it can make sense of much of our existing practice when responsibilities have to be distributed without resort to artificial devices. Consider, for example, a group of hikers out on a trip in the mountains together, where one of the party falls and injures herself. Here we simply take it for granted that the responsibility for bringing aid to the injured member falls in the first place collectively on the whole group, rather than on, say, other climbers who happen to be in the vicinity at the time. By forming ourselves into a group to make the expedition, we create the kind of relationship that generates special responsibilities as a matter of course, and there is no need to invoke a hypothetical contract among the members, or to suppose that somehow members of such a group are better placed to understand their fellow-

¹⁴Unfortunately there is not sufficient space here to consider a close cousin of the capacity principle, the vulnerability principle defended in R. E. Goodin, *Protecting the Vulnerable: A Reanalysis of Our Social Responsibilities* (Chicago: University of Chicago Press, 1985). Goodin's central proposal is that protective responsibilities should be assigned in proportion to vulnerabilities: the more vulnerable P is to A, the greater A's responsibility to protect P from harm. I see this as a cousin of the capacity principle since, in general, the greater an agent's capacity to act, the greater her potential impact on others, and therefore the more vulnerable those others are to her decisions. However P's vulnerability, considered simply by itself, seems to correlate with a responsibility on A's part to avoid causing harm to P, rather than a wider remedial responsibility to rectify P's suffering or deprivation. Suppose that P is the young author of a book on a topic on which I am recognised to be the leading authority, and I am asked to review the book. P is highly vulnerable to my actions: a damaging review will blight his career, a favourable one may launch it. I have a responsibility here not to harm P wilfully or carelessly. If I think the book is a bad one I should say so, but I should take care, for instance, not to indulge my prejudices, precisely because the author is relatively junior, and therefore more vulnerable to my judgement than an established figure would be. So I have a limited responsibility to avoid causing P harm. But I seem to have no special responsibility to protect P from the harm that others may inflict—other reviewers less scrupulous than myself, for instance—and in the event that P is damaged, whether by my own (fair) review, or by other hands, I bear no remedial responsibilities. (If I review the book unfairly, and P's career suffers as a result, then I do have remedial responsibilities, but these are better explained by the moral responsibility principle discussed earlier than by the vulnerability principle; it is the faultiness of my conduct, not the simple fact of P's vulnerability, that creates such responsibilities.)

members' needs than non-members would be, in order to reach the conclusion that when one person gets injured, it is the group that bears the primary responsibility to remedy her condition.¹⁵ But the communitarian principle does less well on two other counts.

First, and only too obviously, it cannot explain why remedial responsibilities sometimes exist in the absence of special relationships of the kind outlined above. But there are at least two cases in which their existence seems pretty clear-cut. One is the case in which A simply injures P without justification, particularly perhaps when the injury involves a violation of P's basic rights (to bodily integrity, say, or to subsistence). That A bears a remedial responsibility in such a case seems uncontroversial, and this is regardless of whether A and P are linked by some kind of communal bond. If I injure a complete stranger, someone not connected to me by nationality, religion or any of the possible forms of communitarian tie, I still owe him a remedy for the injury I have caused.

Likewise, and only a little more controversially, A may be remedially responsible for P simply because he happens to be the only person currently in a position to do anything about P's condition. This assumes of course that P's condition is one of significant deprivation or injury. But if that is the case, then the general duty we all have to aid people who are seriously injured or deprived will devolve upon A simply because he is the only person able to discharge it, either because, for instance, he is the sole passer-by when the skater on the pond falls through the ice, or because he alone has the know-how to sort out P's condition. Once again these responsibilities appear to transcend any communal bonds that might exist between A and P.¹⁶

The second limitation of the communitarian principle is that it has nothing directly to say about how responsibilities are to be distributed *within* the community. It can accommodate the fact that one community may nest inside another, so that A can bear special responsibilities to everyone in C, but also more extensive responsibilities still to everyone in C¹, a sub-set of C. But it cannot go beyond this in distributing responsibilities within C except by importing considerations that are not themselves of a communitarian character. Our belief that in some instances people with greater capacity to help P should bear more of the responsibility for doing so, as suggested by the capacity principle, cannot be generated from within the communitarian perspective alone. And since a full account of remedial responsibilities should aim to get beyond the

¹⁵The best non-reductionist account of such special responsibilities known to me is to be found in S. Scheffler, 'Relationships and responsibilities', *Philosophy and Public Affairs*, 26 (1997), 189–209.

¹⁶One may of course think that where communal bonds also exist this *strengthens* A's obligation, so that if A has to decide which of two endangered skaters to rescue first, he should give precedence to the one who lives in his neighbourhood or belongs to his church; similarly a medical specialist may give priority to treating critically ill patients in his own nation while also acknowledging some responsibility towards similar patients elsewhere.

group level and attribute responsibility to individuals, the communitarian principle taken by itself appears inadequate.¹⁷

II. TAKING STOCK

We seem to have reached an impasse. We have looked at four principles—causal responsibility, moral responsibility, capacity and community—each of which seems *prima facie* plausible as a way of allocating remedial responsibilities, but none of which, on closer inspection, seems adequate by itself. So how should we proceed? There seem to be three main options. The first is to abandon the search for a general theory of remedial responsibilities. The grounds for doing this would be that my initial formulation of the problem—here is a group of deprived or suffering people; whose responsibility is it to put the situation right?—misleadingly amalgamates a range of quite different cases, in each of which a different principle applies. For example, that formulation glosses over the difference between the case in which P's deprivation or suffering is the result of human action, and the case where it has purely natural causes. So instead of looking for a general theory to explain our remedial responsibilities, we need to disaggregate the cases first, and then perhaps construct specific accounts of responsibility to fit each of them in turn.

The second option is to mount a defence of one of the principles canvassed above, and then argue that all of our considered judgements about remedial responsibility can in fact be captured by that principle. This may involve conceding that there are cases in which people are deprived or suffering but in which no-one has a responsibility to help them. For instance, one might argue that agents have remedial responsibilities only when they are morally responsible for the deprivation and suffering in question. Someone who took that line would have to reject my original formulation in a different way: they would have to deny the premise that it is morally intolerable if (remediable) suffering and deprivation are allowed to continue, in other words that where they exist we are morally bound to hold *somebody* (some person or some collective agent) responsible for relieving them.

The third option is to construct a multi-principle theory that combines the four principles I have identified in some fashion. There are various ways in which this might be done, but broadly speaking we can distinguish theories that rank the principles in a certain way—recommend that they should be applied sequentially,

¹⁷In some cases, of course, responsibility does just rest with the group as a whole until it is assigned by some mechanism to individuals, as in the case of the mountaineering group who share responsibility equally for their injured comrade until, say, they draw straws to decide who should go back and call out the Mountain Rescue team. There is nothing to distinguish one member from another in the assigning of responsibility. But this will not be true in general.

for instance—and theories that are more straightforwardly pluralist, in the sense that they ask us to balance or weigh the various principles against one another when responsibilities have to be assigned.

The first option is really a counsel of despair. If our aim in developing a theory of remedial responsibility is eventually to be able to pin these responsibilities on to particular agents and then exert pressure on them to discharge their obligations, a unified theory is a much better tool than a string of sub-theories. But more importantly, few cases in real life will fall neatly into just one of the categories we might construct. The example with which I began illustrates this. When we think about who bears remedial responsibility for Iraqi children, we think about what certain agents have done, what other agents have failed to do, we think about who can claim to have acted with justification, and we also think about who is now best placed to help the children. It is not a simple case in which an identifiable A has unjustifiably injured an identifiable P, nor is it a simple case in which P's condition has purely natural causes. The same will hold in all but the most primitive cases of human deprivation and suffering. So it would be premature to abandon the search for a general theory of remedial responsibility.

The second option involves attempting to bring everything that we want to say about remedial responsibility under the auspices of one of the principles identified above. How might a theory of this kind be constructed? The most plausible candidates, I believe, are a backward-looking theory of moral responsibility and a forward-looking theory, some variant of the capacity view, which focuses on the effective relief of harm and deprivation. A theory of the first kind would hold that agents are only remedially responsible for situations when they are blameable for bringing them about, and a proponent of such a theory might be willing to accept that in cases where no such agent can be found, no remedial responsibilities exist (thus no-one should be held remedially responsible for the victims of natural disasters, though no doubt it would be a morally worthy act to help them). Such a theory looks entirely to the past; the only question it asks is 'who has done this thing?'

A wholly backward-looking theory of this kind finds itself trapped in the following dilemma: either it leaves victims intolerably exposed, in the sense that many injuries will go unremedied, or inadequately remedied, or else it imposes an intolerable burden on agents—this could mean any one of us—by making us responsible for potentially enormous costs resulting from our actions. To see this, consider a case in which some slight act of carelessness on my part results in injuries to Smith that require millions of pounds to compensate. In this instance, we can either narrow the concept of moral responsibility so that I am no longer responsible for everything that results from my careless action—in which case Smith's condition will largely go unremedied—or else we broaden it so that I become remedially responsible for Smith's condition, in which case we would seriously limit our freedom as actors. Every time we acted, we would

face a small but significant risk of landing ourselves with crippling remedial responsibilities.¹⁸ A purely backward-looking theory cannot satisfy our underlying interest in protection against harm and deprivation without burdening us with liabilities that would make everyday life (driving to work, walking down the street) a potentially hazardous business.

Forward-looking theories tell us to assign responsibility in whatever way will best achieve our aim of relieving victims, and thus to ignore the past except insofar as it tells us things about agents that are now relevant in deciding who is most appropriately placed to remedy P's plight. Such theories look at the issue from the side of the victim, and their corresponding weakness is that they assign no intrinsic weight to the value we attach to moral responsibility. Suppose A injures P in a simple and straightforward way. A forward-looking theory asks 'Who should now remedy P's condition?' and A will come into the frame only if we judge it most useful or beneficial to hold him responsible, either in the particular case, or because we think it a beneficial rule to assign remedial responsibility in such cases to the agent who brought about P's condition. But this does not take the past seriously enough: A has done this thing, he can now make amends, so why should we look any further in assigning responsibility? The roundabout reasoning that a forward-looking theory requires seems to violate a basic belief that, at least in simple cases, people should be held responsible for the harm that they do.¹⁹

It seems, therefore, that an acceptable theory of remedial responsibility must make room in some fashion for each of the principles identified above. But how is this to be done? Again, there are different ways of constructing a multi-principle theory. Let us consider some of the more plausible alternatives.

One possibility is that the relevant principles should be applied in sequence: we look to see whether there is any agent who can be held responsible for remedying P's condition under principle X, and if the answer is Yes that settles the matter; if the answer is No we proceed to principle Y, and so forth. The plausibility of such a theory will depend on the sequence chosen, so rather than review all of the variations let me focus on one likely candidate. This tells us to begin by applying the principle of moral responsibility. If we can find an agent who is responsible for P's plight, and also has the capacity to remedy it, then we should hold that agent remedially responsible for P. Failing that—if no morally responsible agent can be identified, or if the agent or agents who bear moral responsibility turn out

¹⁸Insurance provides a mechanism for offsetting this objection, and, as Honoré points out in relation to tort law 'some form of insurance is essential if a system of corrective justice is to operate fairly in modern conditions': T. Honoré, 'The morality of tort law', *Philosophical Foundations of Tort Law*, ed. D. G. Owen (Oxford: Clarendon Press, 1995), p. 90. But it is of course only contingent that such a mechanism exists—so a ground-level theory of remedial responsibility cannot presuppose that it does (that is, we cannot argue that it is essentially fair to hold people remedially responsible for compensating all of the harm that may eventuate from their actions on the grounds that, if insurance mechanisms exist, they can choose to protect themselves from incurring excessive burdens).

¹⁹Some considerations in support of this basic belief are advanced in Honoré, 'Responsibility and luck'.

to be incapable of supplying the remedy—we invoke another principle—causal responsibility, say, or community. Clearly there are different ways of completing the theory. But we need go no further, because I do not think that the first move is defensible.

The problem is that moral responsibility is a matter of degree, and degrees of moral responsibility for P's condition do not necessarily correlate with other relevant features, especially with the capacity to relieve P's suffering or deprivation. We can lose sight of this fact by thinking in terms of very simple cases, namely those in which A is solely responsible for injuring P, does so by virtue of a deliberate act, and as a result has the resources which could now be used to compensate for the injury. But although such cases undoubtedly occur, we are more often confronted with ones in which A bears a lesser degree of moral responsibility, either by virtue of the character of his action—he acted negligently, for instance, rather than deliberately, or he acted with justification, harming P in pursuit of some greater good—or because responsibility has to be shared between a number of different agents. Moreover the fact that A played some part in bringing about P's condition does not entail that he derived any tangible benefit from doing so. So if we think about a case such as the Iraqi children, we might well conclude that each of the agents identified in the case—the United Nations, Saddam's clique, and the Iraqi people—bore some share of moral responsibility for the suffering of those children, though the nature of the responsibility would differ in each case. The United Nations might be accused of pursuing a justifiable end by unacceptable means, the suffering of the children being a foreseeable side-effect of the sanctions policy. Saddam might be accused, more harshly, of deliberately allowing the suffering to occur for propaganda purposes. About the Iraqi people, it could be said that their causal role in bringing about the suffering was almost entirely negative, but that in the circumstances they could reasonably have been expected to oppose Saddam's regime more effectively. If those judgements are accepted, we would have to conclude, first, that no agent can be singled out as uniquely morally responsible for the situation we want to have remedied; and, second, that the agent bearing the largest share of responsibility (Saddam's clique) is also worst placed to bring the children the help that they need. It is probably wrong to say that the Iraqi government is incapable of finding the resources in question. But clearly it would be very much easier for the Western powers to do so. In these circumstances, is it right to let our judgements of remedial responsibility be determined entirely by our prior beliefs about how far different agents are morally responsible for creating the situation that needs a remedy?

This suggests that, rather than applying principles of responsibility in strict sequence, our approach should be more openly pluralist: we should simply look to see which principle or principles apply in a particular case, and if we find that more than one applies, we should weigh their respective strengths. But before

reaching this conclusion, we should consider a second way of ordering the principles. This approach sees them not as competing with one another, but as addressing different aspects of the allocation of responsibility. Again I shall focus on one particular version of this approach.

The position I have in mind holds that we need to distinguish *immediate* responsibility for relieving harm and suffering from *final* responsibility.²⁰ Where people are in distress or in danger of further injury, we need to identify the agents best placed to help them in the short term. But these may not be the agents who should bear the costs of such action in the long term. So A may be immediately responsible for relieving P, because A is the agent in a position to offer aid directly, but final responsibility may be B's, in which case B may have to compensate A for the resources she has provided to B or for incidental costs arising from the relief effort.

Adopting such a view, we might conclude that *capacity*, and to some extent *community*, are relevant principles when immediate responsibilities are being distributed, because these are criteria that tell us who is best able to relieve P's condition quickly and effectively. *Moral responsibility*, *causal responsibility*, and perhaps *community* again, are invoked when final responsibilities are the issue.²¹ Now clearly there are cases where something like that picture seems to apply. If I negligently allow my child to fall into the river, but I cannot swim, then immediate responsibility for the rescue may fall on a passer-by who can, while I remain ultimately responsible for my child's welfare once he is pulled out, and for any damage suffered by the rescuer. But what makes the picture appropriate is precisely the immediacy of the harm: if the child is not saved at once he will drown. It is that feature that makes capacity the overriding consideration in the short term. In many other cases, however, the deprivation or suffering that motivates our assigning of responsibility is relatively stable, in the sense that we have no reason to expect the situation to deteriorate suddenly, and here it seems that capacity becomes one relevant consideration to consider, alongside the others. And even where harm is immediate, capacity may not trump the other principles in a straightforward way: where several swimmers could rescue the drowning child, we may not simply pick the strongest, but, for instance, the person who caused him to fall in, or the negligent parent who should have taken more care to keep him away from the water's edge.

III. THE CONNECTION THEORY

It appears that we must settle for a pluralist approach to distributing responsibilities. Attempts to eliminate all but one of the principles we have

²⁰I am grateful to an anonymous referee for this *Journal* for suggesting this position.

²¹Community might appear in both places because, on the one hand, it can serve to identify agents who are physically proximate to P and therefore likely to be well placed to help him, and on the other hand, it picks out a group of agents who have a special concern for P's welfare.

unearthed lead to unacceptable results. Attempts to impose a fixed order of application on those principles also lead us astray. But can a pluralist approach give us a satisfying theory of remedial responsibility? I shall sketch such a theory, which I propose to call the *connection* theory.²²

This begins with the observation that all of us have a strong interest in the existence of mechanisms that protect us from harm and injury; more particularly, in the event that we find ourselves in a bad condition which it is difficult or impossible for us to remedy through our own devices, we want there to be some way of assigning responsibilities such that an identifiable A (or perhaps A, B and C taken together) becomes responsible for rectifying our situation. We want A to feel that he is responsible, and to act accordingly, and we want everyone else to make the same judgement and therefore to put pressure of various kinds on A if he fails to act. Responsibility that is widely dispersed is no good, because then everyone will attempt to hang back in the hope that someone else will step in first, no-one will be particularly liable to censure if the bad condition is not remedied, and so on.

The interest of potential P's in having clearly defined responsibilities is evident, reflecting the fact that all of us are vulnerable to outside events that may seriously harm our interests or threaten our lives. But equally, it may not matter so much, from P's point of view, which particular A is held responsible. There may be many agents who are able to remedy P's condition; it matters considerably to P that one such agent, or one group of agents, should be singled out, for the reasons given above, but apart from that P may be indifferent which agent this happens to be. In these circumstances we will fix responsibility on the agent who is already connected to P in some way; if several agents are so connected, we will choose the one whose link to P is strongest, or else, depending on the circumstances, divide up responsibility according to relative strength of connection. (Whether one agent is singled out, or responsibility shared between several, may depend on whether P's condition is better remedied by a single agent taking action or several agents acting in concert.)

What kind of connection is relevant here? My proposal is that we should return to the principles explored in the second section of the article, and now see them as specifying forms of connection between A and P that may, in particular cases, be sufficient to establish A's responsibility for remedying P's condition. That is, A may be remedially responsible for P either because he is causally responsible for P's condition, or morally responsible for it, or has the (special)

²²In thinking about the general shape of this theory, I have drawn inspiration from Hume's theory of justice and property, as set out particularly in *A Treatise of Human Nature*, ed. L. A. Selby-Bigge, 2nd edn revised P. H. Niddich (Oxford: Clarendon Press, 1978), Pt III, secs 2–3, 6. Hume argued that we have a shared strong interest in the existence of rules that stably assign possessions to persons as their property. The specific rules employed, however, depended less on utilitarian considerations than on common mental dispositions that lead us to connect persons to material objects in particular ways. The connection theory of remedial responsibility proposed here has a similar general structure, though as readers have pointed out there are also significant disanalogies between the two theories.

capacity to rectify it, or already has a communal relationship with P. Any of these relations—causal responsibility, moral responsibility, capacity, or community—may establish the kind of special link between A and P that enables us to single out A as the one who bears the responsibility for supplying the resources that will remedy P's condition.

In many cases there will be independent moral reasons for using a particular connection as a basis for assigning remedial responsibility. This applies most obviously where the connection consists in A's moral responsibility for P's injury. There are two independent reasons for holding A remedially responsible in this case: first, where A has unjustly benefited from the injury he has inflicted on P—he has stolen something of P's, or exploited him, for example—then if A is made to compensate P by returning what he has taken or in some other way undoing the damage he has inflicted, then this will help to cancel out A's unjust gain, and so restore justice between them. Second, even if A has not benefited from his actions, he has wronged P, and should therefore make recompense to P as a way of acknowledging the wrong he has committed; remedying P's injury is an obvious way to do this. In the cases of capacity and community, too, we can provide a plausible rationale for basing remedial responsibilities on these forms of connection (these rationales will be quite different from the one just given in the case of moral responsibility). But causal responsibility, in the absence of the other forms of connection, seems not to have its own ethical rationale. If A is causally responsible for P's condition, but he is neither morally responsible for it (say he could not possibly have foreseen the results of his action) nor especially well placed to assist P, nor linked to P by special communal ties, there seems no independent reason to hold him remedially responsible for sorting P out. Yet we do seem to think that a bare causal connection is enough to generate special responsibilities, as the case of the innocent pedestrian who knocks down another shows. Perhaps it might be argued that in such cases holding the causal agent responsible will create incentives to take special care not to inflict accidental injury; thus the rationale is indirectly utilitarian.²³ But this interpretation seems forced: if the link between action and outcome is such that the agent could not be expected to anticipate the injury to P, then how is A supposed to alter his behaviour? I suggest instead that causal responsibility as a source of remedial responsibility confirms the theory I am putting forward: we need some way of identifying an A to hold remedially responsible for P's condition, so in the

²³Another suggestion might be that causal responsibility always carries traces of moral responsibility with it. How can we be sure that the pedestrian who was apparently taking good care to watch his step did not in fact have a momentary lapse of concentration and fail to notice the raised paving stone which he ought to have noticed? Or alternatively perhaps we are all to some small degree morally responsible for *everything* that results from our action, no matter how remote the causal chain, or how careful we have been. In ways such as these, we can put a moral gloss on causal responsibility. But then again, might not this be a case of the tail wagging the dog? Might it not be that *because* we want to hold causal agents remedially responsible for the harm they bring about, in certain cases, we are driven to impute moral responsibility to them, in defiance of our more usual (and defensible) understanding of that idea?

absence of any other link between potential A's and P we fix upon the purely physical link of causality. This, admittedly, can be trumped fairly readily once other forms of connection come into play; but the fact that we are prepared to rely on bare causal connection in the first place underlines the necessity of finding some non-arbitrary way of assigning responsibility to a particular agent.

The connection theory successfully offers practical guidance in cases where only one of its four constituent principles is satisfied. There is nothing paradoxical, according to the theory, in assigning remedial responsibility on the basis of causation, say, in one case and on the basis of communitarian relations in the next case. Our overriding interest is to identify an agent who can remedy the deprivation or suffering that concerns us, and in pursuit of that aim we fix on whoever is linked to P according to one of the theory's four criteria, about which there is widespread agreement. Where two or more of the principles apply, the theory tells us to look at the strength of the various connections. Thus if A is weakly linked to P by virtue of moral responsibility, whereas B is strongly linked to P by virtue of capacity (B is in a far better position to remedy P's condition than any other agent), the theory instructs us to hold B remedially responsible. In some cases it may recommend dividing responsibility between two or more agents, where this makes practical sense, and the ties are of comparable strength. On the connection theory there is no algorithm that can tell us to apply principle 1 first and then move on to principle 2, and so forth—my reasons for rejecting such an algorithm have already been given. This means, of course, that when connections have to be weighed against each other, we can do no more than appeal to shared moral intuitions about which is the stronger.

The strength of the connection theory, as I see it, is that it treats the obligation to relieve deprivation and suffering as of overriding concern. By using multiple criteria, we ensure that there is always *some* agent who can be assigned responsibility for remedying P's condition. At the same time, it makes room for other moral considerations, such as the deeply held belief that where we can point the finger at a particular A as being morally responsible for the harm suffered by P, it is A himself who should remedy the harm wherever possible. Single-principle theories, I have suggested, will inevitably run up against such beliefs sooner or later. And although the connection theory is internally complex, this complexity may simply mirror the complexity of real-world cases in which remedial responsibility has to be assigned. If there were a simple answer to questions such as who is responsible for the current plight of Iraqi children, we would not argue about it politically in the way that we do. The connection theory does not offer a mechanical answer to questions of that kind, but it provides a way of thinking about them—highlighting their complexity—that may in the end prove to be more illuminating.