



Philosophical Review

Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing

Author(s): Warren S. Quinn

Source: *The Philosophical Review*, Vol. 98, No. 3 (Jul., 1989), pp. 287-312

Published by: Duke University Press on behalf of Philosophical Review

Stable URL: <http://www.jstor.org/stable/2185021>

Accessed: 24/10/2009 11:14

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=duke>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Duke University Press and *Philosophical Review* are collaborating with JSTOR to digitize, preserve and extend access to *The Philosophical Review*.

<http://www.jstor.org>

ACTIONS, INTENTIONS, AND CONSEQUENCES:
THE DOCTRINE OF DOING AND ALLOWING¹

Warren S. Quinn

Sometimes we cannot benefit one person without harming, or failing to help, another; and where the cost to the other would be serious—where, for example, he would die—a substantial moral question is raised: would the benefit justify the harm? Some moralists would answer this question by balancing the good against the evil. But others deny that consequences are the only things of moral relevance. To them it also matters whether the harm comes from action, for example, from killing someone, or from inaction, for example, from not saving someone. They hold that for some good ends we might properly allow a certain evil to befall someone, even though we could not actively bring that evil about. Some people also see moral significance in the distinction between what we intend as a means or an end and what we merely foresee will result incidentally from our choice. They hold that in some situations we might properly bring about a certain evil if it were merely foreseen but not if it were intended.

Those who find these distinctions morally relevant think that a benefit sufficient to justify harmful choices of one sort may fail to justify choices no more harmful, but of the other sort.² In the case of the distinction between the intentional and the merely foreseen, this view is central to what is usually called the Doctrine of Double Effect (DDE). In the case of the distinction between action and inaction, the view has no common name, so for convenience we may call it the Doctrine of Doing and Allowing (DDA). (Because harm resulting from intentional inaction has, typically, been allowed to occur.) Absolutist forms of either doctrine would simply

¹Thanks to Rogers Albritton, Tyler Burge, Philippa Foot, Matthew Hanser, Thomas Nagel, Michael Thomson, Derek Parfit, T. M. Scanlon, and to the editors of *The Philosophical Review* for valuable suggestions and criticisms.

²Harm here is meant to include any evil that can be the upshot of choice, for example, the loss of privacy, property, or control. But to keep matters simple, my examples will generally involve physical harm, and the harm in question will generally be death.

rule out certain choices (for example, murder or torture) no matter what might be gained from them. Nonabsolutist forms would simply demand more offsetting benefit as a minimum justification for choices of one sort than for equally harmful choices of the other sort.

In this paper I shall examine the Doctrine of Doing and Allowing.³ My aim is twofold: first, to find the formulation of the distinction that best fits our moral intuitions and second, to find a theoretical rationale for thinking the distinction, and the intuitions, morally significant. Both tasks are difficult, but the former will prove especially complex. What we find in the historical and contemporary literature on this topic is not a single clearly drawn distinction, but several rather different distinctions conforming roughly but not exactly to the distinction between what someone does and what he does not do. Special cases of inaction may be treated by an author as belonging, morally speaking, with the doings, and special cases of doing as belonging with the inactions. So in searching for the proper intuitive fit, we shall have to be alert to the possibility that the distinction between action and inaction (or between doing and allowing) is only a first approximation to the distinction we really want.

In evaluating various formulations of the doctrine I shall need special test cases. These will often involve improbable scenarios and repetitive structural elements. This is likely to try the reader's patience (he or she may begin to wonder, for example, whether we are discussing the morality of public transportation). But it may help to recall that such artificialities can hardly be avoided anywhere in philosophy.⁴ As in science, the odd sharp focus of the test cases is perfectly compatible with the general importance of the ideas being tested. And the DDA is, I think, of the greatest general significance, both because it enters as a strand into many real moral issues and because it stands in apparent opposition to that most general of all moral theories, consequentialism.

Before beginning, I should emphasize that both the DDE and, especially, the DDA apply more directly to moral justification than to other forms of moral evaluation. It is therefore open to a de-

³I shall examine the DDE in a subsequent paper, "Actions, Intentions, and Consequences: The Doctrine of Double Effect."

⁴Think of Gettier cases, brain transplants, teletransporters, etc.

fender of the DDA to admit that two *unjustified* choices that cause the same degree of harm are equally *bad*, even though one choice is to harm and the other not to save. I note this only because some writers have looked for such pairs in hope of refuting the doctrine.⁵ Take the well-known example of an adult who deliberately lets a child drown in order to inherit a family fortune.⁶ The act seems so wicked that we understand the point of saying that it is no better than drowning the child. But if so, how can we hold that the difference between killing and letting die matters morally?

This objection seems to presuppose that if letting someone die is ever more acceptable, *ceteris paribus*, than killing someone, it must be because some intrinsic moral disvalue attaches to killing but not to letting die. And if so, this intrinsic difference must show up in all such cases.⁷ But the doctrine may, and I shall argue should, be understood in a quite different way. The basic thing is not that killing is intrinsically worse than letting die, or more generally that harming is worse than failing to save from harm, but that these different choices run up against different kinds of rights—one of which is stronger than the other in the sense that it is less easily defeated. But its greater strength in this sense does not entail that its *violation* need be noticeably worse.

Such relations between rights are possible because moral blame for the violation of a right depends very much more on motive and

⁵For example, Michael Tooley in "Abortion and Infanticide," *Philosophy and Public Affairs* 2 (1972), p. 59.

⁶From James Rachels, "Active and Passive Euthanasia," *The New England Journal of Medicine* 292 (1975), p. 79.

⁷In "Harming, Not Aiding, and Positive Rights," *Philosophy and Public Affairs* 15 (1986), pp. 5–11, Frances Kamm rightly makes us distinguish between two ways in which killing might be intrinsically worse than letting die: a) killing might have some bad essential feature that cannot attach to letting die or b) killing might have some bad essential feature that, while not essential to letting die, can nevertheless be present in cases of letting die. If (b) is true then the moral equivalence of the two cases in which the child drowns would not establish a general moral equivalence between killing and letting die. For letting the child drown might be a special case in which letting die has the bad feature essential to killings but not to lettings die. And even apart from Kamm's point, the idea that intrinsically nonequivalent parts must always make an overall evaluative difference when embedded in identical contexts seems wrong. Consider aesthetics. There may certainly be important intrinsic aesthetic differences between two lampshades even though they create an equally bad overall impression when placed on a certain lamp.

expected harm than on the degree to which the right is defeasible. Your right of privacy that the police not enter your home without permission, for example, is more easily defeated than your right that I, an ordinary citizen, not do so. But it seems morally no better, and perhaps even worse, for the police to violate this right than for me to. So there is nothing absurd in saying that the adult acts as badly when he lets the child drown as when he drowns the child, while insisting that there are contexts in which the child would retain the right not to be killed but not the right to be saved.

I.

The Doctrine of Doing and Allowing has been most notably defended in recent moral philosophy by Philippa Foot.⁸ It will be convenient, therefore, to begin with two of the examples she uses to show the intuitive force of the doctrine.⁹ In Rescue I, we can save either five people in danger of drowning at one place or a single person in danger of drowning somewhere else. We cannot save all six. In Rescue II, we can save the five only by driving over and thereby killing someone who (for an unspecified reason) is trapped on the road. If we do not undertake the rescue, the trapped person can later be freed. In Rescue I, we seem perfectly justified in proceeding to save the five even though we thereby fail to save the one. In Rescue II, however, it is far from obvious that we may proceed. The doctrine is meant to capture and explain pairs of cases like these in which consequential considerations are apparently held constant (for example, five lives versus one) but in which we are inclined to sharply divergent moral verdicts.

⁸In "The Problem of Abortion and the Doctrine of the Double Effect," in *Virtues and Vices and Other Essays* (Berkeley, Calif.: University of California Press, 1978), pp. 19–32, Foot argued that the distinction between doing and allowing could do all the work usually credited to the distinction between the intentional and the merely foreseen. In "Killing and Letting Die," Jay Garfield, ed., *Abortion: Moral and Legal Perspectives* (Amherst, Mass.: University of Massachusetts Press, 1984), pp. 178–185 and even later in "Morality, Action and Outcome," in Ted Honderich, ed., *Morality and Objectivity* (London, England: Routledge and Kegan Paul, 1985), pp. 23–38, she withdraws this claim, arguing instead that any intuitively adequate morality must assign an independent moral significance to the distinction between doing and allowing.

⁹From "Killing and Letting Die," p. 179.

The first order of business is to get clearer on the crucial distinction that the doctrine invokes. In effect, the DDA discriminates between two kinds of agency in which harm comes to somebody. It discriminates *in favor of* one kind of agency (for example, letting someone drown in Rescue I) and it discriminates *against* the other kind (for example, running over someone in Rescue II).¹⁰ That is, it makes these discriminations in the sense of allowing that the pursuit of certain goods can justify the first kind of harmful agency but not the second. I shall call the favored kind of agency *negative*, since on any plausible account it is usually a matter of what the agent does *not* do. For parallel reasons, I shall call the disfavored kind of agency *positive*. But, as indicated earlier, the distinction between positive and negative agency may or may not line up exactly with the ordinary distinction between doing and allowing or action and inaction. We may discover, as we consider various special circumstances, that certain actions function morally as allowings and certain inactions as doings. So let us begin by sifting various proposals for spelling out the nonmoral difference between the two kinds of agency.

One such proposal comes from some brief passages in the *Summa Theologiae* where Aquinas could be taken to suggest that the difference between the two forms of agency is one of voluntariness.¹¹ In harmful positive agency, the harm proceeds from the will of the agent while in harmful negative agency it does not.¹² St. Thomas seems to think that foreseeable harm that comes from action is automatically voluntary. But he thinks that foreseeable

¹⁰It seems clear than an agent's *not* doing something (for example, not saving someone from drowning) can be morally evaluated as justified, unjustified, right, or wrong, in precisely the sense in which these terms apply to actions. I shall therefore speak of assessing the justification or the rightness of someone's *agency* in some matter, meaning by this an evaluation of his knowingly acting or not acting.

¹¹*Summa Theologiae* XVII (Cambridge, England: Blackfriars, 1970), 1a2ae Q. 6 article 3, pp. 15–16. The terms “positive agency” and “negative agency” are not, of course, St. Thomas's. This is the interpretation that he might give to them.

¹²In speaking of an inaction as harmful or as producing harm (or in speaking of harm as coming from it) I am not begging the question against Aquinas. For I mean these expressions only in the weak sense of connecting the inaction with a harmful upshot, and not in any sense that would imply that the harm was voluntary.

harm coming from inaction is voluntary only when the agent could and *should* have acted to prevent it. Positive agency would therefore include all foreseeably harmful actions and those foreseeably harmful inactions that could and should have been avoided. And negative agency would include the foreseeably harmful inactions that could not or need not have been avoided.

But what kind of “should” (or “need”) is this? If we take it to be moral, the doctrine becomes circular.¹³ Inactions falling under positive agency are harder to justify than inactions falling under negative agency. Why? Because by definition the latter need not have been avoided while the former, if possible, should have been.

We could, however, avoid the circularity by taking the “should” to be premoral, reflecting social and legal conventions that assign various tasks to different persons. And we might think that these conventions play a central role in an important premoral, but morally relevant, notion of causality.¹⁴ The helmsman’s job is to steer the ship, and this is why we say that it foundered *because* of his careless inaction. The loss of the ship would thus be like the death in Rescue II, which happens because of what we do. And both cases would contrast with the death in Rescue I, which we do not, in the relevant sense, bring about. Voluntariness would thus be seen as a distinctive kind of causal relation linking agency and its harmful upshots in the cases of action and conventionally proscribed inaction (positive agency), but not in the case of conventionally permitted inaction (negative agency). So formulated, the doctrine would not only be clear but would have an obvious rationale. Harmful negative agency is easier to justify because in such cases the harm cannot, in the relevant causal sense, be laid at the agent’s door.

I have two objections to this proposal. First, there is little reason to treat most instances of the neglect of conventional duty as posi-

¹³It may be, of course, that Aquinas’s account of the voluntary is not meant as part of the theory of justification and is therefore not directed to the distinction between positive and negative agency. It might instead be part of the theory of praiseworthiness and blameworthiness, which presupposes an independent account of what can and cannot be justified. If so, there could be no charge of circularity. For harmful inaction clearly does deserve blame only if it could and should, morally speaking, have been avoided.

¹⁴I am indebted here to Michael Thompson, who thinks that something like this is suggested in the work of Elizabeth Anscombe.

tive agency. We can usually explain in other ways just why morality takes these tasks so seriously. If human communities are to thrive, people will have to perform their social roles. That is why, in a variant of Rescue I, the private lifeguard of the lone individual might not be morally permitted to go off to save the five, even though a mere bystander would be. To explain this difference, we would not also need to suppose that the private lifeguard and bystander stand in different causal relations to the person's death. There is, moreover, room to think that the special duty of the private lifeguard should be put aside, especially if his employer is a pampered rich man, and the five are too poor to afford personal attendants. But this kind of circumstance would have no justificatory force where death was the upshot of clearly positive agency. In Rescue II, for example, it would not matter that the man trapped on the road was rich and spoiled while the five were poor and worthy.

My second objection is more general. The type of proposal we are examining relationalizes the special moral opprobrium attaching to positive agency by reference to its special causal properties. Since negative agency is not, in the intended sense, the cause of its unfortunate upshots, the moral barriers against it are lower. But this leaves the doctrine open to a serious criticism. For there are other conceptions of causality according to which we are in (the original) Rescue I every bit as much a cause of death as in Rescue II. What matters, according to these conceptions, is whether a nonoccurrence necessary for a given effect was, relative to a certain standard background, surprising or noteworthy. In this sense, we may say that a building burns down because its sprinklers failed to work, even though their failure was traceable to the diversion of water to another more important fire. That the diversion was quite proper is nothing against the claim that the failure of the sprinklers helped cause the loss of the building. And something similar holds for Rescue I. The fact that we did not save the one because we quite properly saved the others would not show that his death was not in part due to our choice.

So even if there is a causal notion that corresponds to Aquinas's idea of the voluntary, it is in competition with other causal notions that may seem better to capture what is empirically important in scientific and ordinary explanation. And it is arguable that the defense of the doctrine should not depend on a causal conception

that we would otherwise do without. If the doctrine is sound it ought to remain plausible on an independently plausible theory of causation. In any case, this is what I shall assume here. So I shall grant opponents of the doctrine that the permissible inactions we are considering, no less than the impermissible actions, are partial causes of their harmful upshots. This will force me to try to make sense of the doctrine on other grounds.

But this still leaves the task of stating the nonmoral content of the distinction between harmful positive and harmful negative agency. Perhaps the difference should, after all, be put in the most simple and straightforward way, as the difference between action that produces harm and inaction that produces harm. If we think of action along the lines proposed by Elizabeth Anscombe and taken up by Donald Davidson—a conception whose basic outline I propose to adopt—individual actions are concrete particulars that may be variously described.¹⁵ To say that John hit Bill yesterday is to say that there was a hitting, done by John to Bill, that occurred yesterday. To say that John did not hit Bill, on the other hand, is to say that there was no such hitting. Taking things this way, the distinction between harmful positive agency and harmful negative agency would be the distinction between harm occurring because of what the agent does (because of the existence of one of his actions) and harm occurring because of what the agent did not do but might have done (because of the noninstantiation of some kind of action that he might have performed).¹⁶

Surprisingly, most moral philosophers who write on these matters reject this way of drawing the distinction. Jonathan Bennett, a severe critic of the DDA, dismisses Davidson's conception of

¹⁵G. M. A. Anscombe, *Intention* (second edition) (Oxford, England: Blackwell, 1963), especially sec. 26, pp. 45–47. See also Donald Davidson, *Essays on Actions and Events* (Oxford, England: Clarendon Press, 1980). See there “The Logical Form of Action Sentences,” pp. 105–122; “Criticism, Comment and Defence,” pp. 122–144, esp. pp. 135–137; and “The Individuation of Events,” pp. 163–180.

¹⁶What I see as right in the Anscombe-Davidson view is the suggested metaphysics—the claim that action is a matter of the presence of something and inaction a matter of its absence. And I think that our intuitions about whether something is an action or inaction as we think about it morally are metaphysically relevant. So I am not greatly worried that someone pursuing the Anscombe-Davidson line might discover criteria of action and inaction that would radically conflict with our judgments in moral thought.

action without argument.¹⁷ Most likely he minds its failure to provide a clear criterion for distinguishing action from inaction in all cases, one that would tell us, for example, whether observing a boycott (by not buying grapes) or snubbing someone (by not acknowledging his greeting) consists in doing something by way of inaction or simply in deliberately not doing something. Bennett is reluctant to assign moral work to any distinction that leaves some cases unclear, especially where there is no theoretically compelling reductionistic theory for the clear cases. But I am disinclined to adopt such a standard. Almost no familiar distinction that applies to real objects is clear in all cases, and theoretical reducibility is a virtue only where things really are reducible. In any case, the imposition of such a standard would shut down moral theory at once, dependent as it is on the as yet unreduced and potentially vague distinctions between what is and is not a person, a promise, an informed consent, etc.

But Bennett is not simply negative. He proposes an ingenious and, for limited applications, clearly drawn distinction between positive and negative *facts* about agency as a respectable way of formulating the doctrine.¹⁸ (Not of course to save it, but to expose it.) Roughly speaking, an event is brought about by someone's positive instrumentality, as Bennett calls it, when the event is explained by a relatively strong fact about the agent's behavior—for example, that he moved in one of a limited number of ways. Negative instrumentality, on the other hand, explains by reference to relatively weak facts about behavior—for example, that the agent moved in any one of a vast number of ways.

The trouble is that this distinction gets certain cases intuitively wrong. Bennett imagines a situation in which if Henry does nothing, just stays where he is, dust will settle and close a tiny electric circuit which will cause something bad—for example, an explosion that will kill Bill.¹⁹ If Henry does nothing, he is by Bennett's criterion positively instrumental in Bill's death. (For only one

¹⁷"Morality and Consequences," *The Tanner Lectures on Human Values II* (Salt Lake City, Utah: University of Utah Press, 1981), pp. 54–55.

¹⁸Bennett, pp. 55–69.

¹⁹Bennett, pp. 66–68. If Henry's body were *activating* the device—if he were depressing a trigger or conducting a current—we might see his agency as positive despite his motionlessness. But Bennett doesn't assign any such role to Henry's body.

of Henry's physical actions, staying still, will cause the death, while indefinitely many will prevent the death.) But suppose Henry could save five only by staying where he is—suppose he is holding a net into which five are falling. Surely he might then properly refuse to move even though it means not saving Bill. For his agency in Bill's death would in that case seem negative, much like that in Rescue I.

Bennett also misses the opposite case. Suppose the device will go off only if Henry makes some move or other. In that case his instrumentality in the death would, for Bennett, be negative. But those who would rule out Rescue II would surely not allow Henry to go to the rescue of five if that meant setting off the device. For his agency in the death of Bill would in that case seem positive.²⁰ Bennett's distinction, however admirable in other ways, is not the one we seem to want. Perhaps this is already clear when we reflect that, according to him, the instrumentality of someone who intentionally moves his body (in, for example, following the command "Move in some way or other—any way you like!") is negative.

Philippa Foot also rejects the idea that the distinction between positive and negative agency is that between action and inaction.²¹ She claims that it would not make any interesting moral difference if respirators (presumably sustaining patients who would otherwise die) had to be turned on again each day. Active turning off and passive not turning on would be morally the same. To be relevant to the present issues, her idea must be that this would not make a difference even in cases where some great good could come about

²⁰That Bill's death would in this case be a side-effect of the rescue does not distinguish it from Rescue II. For in neither case is the death of the one intended. It might be objected that here but not in Rescue II the killing would not be *part* of the rescue. But if Henry's movement sets off the explosion (for example, by triggering a fuse sensitive to movement) then Henry's killing Bill does seem part of the rescue, at least in the sense that he kills Bill by the very movements that form part of the rescue attempt. Of course there could be circumstances in which Henry's movement would not so much set off the explosion as allow it to be set off. Suppose, for example, Henry's remaining where he is prevents dust from settling upon and thereby triggering an explosive device below him. In such a case, I agree that he might go off to save the five. For although he will be active in Bill's death, his agency will involve taking his body from where it would save Bill to where he can make use of it to save the five, a special circumstance that I shall discuss later.

²¹"Morality, Action and Outcome," p. 24.

only if a particular respirator were not running. Let us see whether this is right. Suppose there are temporary electrical problems in a hospital such that the five respirators in Ward B can be kept going only if the one in Ward A is off. On Foot's view it should not matter whether a hospital attendant keeps the five going by shutting down the one or, in case it is the kind that needs to be restarted, by simply not restarting it.

It would be very odd to think that if the single respirator were already off, the attendant would be required to restart it even if that meant shutting down the five in Ward B. So Foot's idea must imply that if the single respirator were running, the attendant could just as properly shut it down to keep the others running. Now while there seems something more objectionable about shutting the respirator down, I think that all things considered it might be permitted. One reason is that we could perhaps see it as a matter of the hospital's allocating something that belongs to it, a special kind of circumstance that we shall consider later. But suppose the hospital is an unusual one in which each patient must provide his own equipment and private nursing care. Suppose further that you are an outsider who happens for some reason to be the only person on the scene when the electrical problem arises. In this case, it seems to matter whether you keep the respirators in Ward B going by not restarting the one in Ward A (it being of the type that needs restarting and the private nurse having failed to show up that day) or whether you actually shut it down. The first case seems rather like Rescue I and the second uncomfortably like Rescue II.

Foot goes on to offer what she takes to be a different and better interpretation of the distinction. She thinks what matters is not the difference between action and inaction but the difference between two relations an agent can have to a sequence of events that leads to harm. It is one thing to *initiate* such a sequence or to *keep it going*, but quite another to *allow it to complete itself* when it is already in train.²² Agency of the first two kinds is positive, while agency that merely allows is negative. One problem with this account arises when we try to explain the difference between allowing a sequence to complete itself and keeping it going when it would otherwise

²²Ibid., p. 24, including footnote 2 on p. 37.

have stopped. We might have thought that the former was a matter of doing nothing to stop it and the latter was a matter of doing something to continue it. But that would seem to take us back to the rejected distinction between action and inaction.

Another problem concerns forms of help and support which do not seem to consist in keeping already existing dangerous sequences at bay. Suppose I have always fired up my aged neighbor's furnace before it runs out of fuel. I haven't promised to do it, but I have always done it and intend to continue. Now suppose that an emergency arises involving five other equally close and needy friends who live far away, and that I can save them only by going off immediately and letting my neighbor freeze. This seems to be more like Rescue I than Rescue II, but it doesn't appear to be a case in which I merely allow an already existing fatal sequence to finish my neighbor off. For he was not already freezing or even, in some familiar sense, in danger of freezing before the emergency arose. Or if we think he was in danger, that danger was partly constituted by what I might fail to do. We might simply stipulate, of course, that any fatal sequence that appears to arise from a *failure* to help someone is really the continuation of a preexisting sequence. But then we seem to be falling back on the notion of inaction as fundamental.

II.

I am therefore inclined to reject Bennett's and Foot's positive suggestions, despite their obvious attractions. May we then return to the simple and straightforward way of drawing the distinction, as between harm that comes from action and harm that comes from inaction? I think not. Cases involving the harmful *action of objects or forces* over which we have certain powers of control seem to demand a more complex treatment. Consider, for example, the following variant of Rescue II (call it Rescue III). We are off by special train to save five who are in imminent danger of death. Every second counts. You have just taken over from the driver, who has left the locomotive to attend to something. Since the train is on automatic control you need do nothing to keep it going. But you can stop it by putting on the brakes. You suddenly see someone trapped ahead on the track. Unless you act he will be killed. But if you do stop, and then free the man, the rescue mission will be aborted. So you let the train continue.

In this case it seems to me that you make the wrong choice. You must stop the train. It might seem at first that this is because you occupy, if only temporarily, the role of driver and have therefore assumed a driver's special responsibility to drive the train safely. But, upon reflection, it would not make much moral difference whether you were actually driving the train or merely had access to its brake. Nor would it much matter whether you were in the train or had happened upon a trackside braking device.²³ The important thing from the standpoint of your agency is that you *can* stop the train and thereby prevent it from killing the one.

But this is not the only thing that matters, as can be seen in a different kind of case. Suppose, in a variant of Rescue I (Rescue IV), you are on a train on which there has just been an explosion. You can stop the train, but that is a complicated business that would take time. So you set it on automatic forward and rush back to the five badly wounded passengers. While attending to them, you learn that a man is trapped far ahead on the track. You must decide whether to return to the cabin to save him or stay with the passengers and save them.

May you stay? I think you may.²⁴ We would be more tolerant of

²³Suppose that you and a friend are off, by car, on a rescue mission that unexpectedly turns into Rescue II. You are sitting in the passenger seat, and your friend is driving. For some reason he hasn't noticed the trapped person, but you have. If you do nothing, your friend will inadvertently run over and kill the man. Can you really think that the end of rescuing the five would *not* justify your friend the driver in deliberately killing the man, but *would* justify you in keeping silent (or in not pulling up the hand brake)? I find this implausible. And it seems equally implausible to suppose that your obligation to yell or pull the brake comes from your having, temporarily, assumed the role of driver. What matters is that the mission has become illicit precisely because, as you can see, it requires that someone be killed. So it has also become illicit to try to *further* the mission, whether by deliberate action or omission.

²⁴At least if you are not the driver or his designated replacement—that is, someone charged with a special moral responsibility to see to it that the train kills no one. If you have that responsibility but lack a special duty toward the injured people (you are not also their doctor), then there would be something extra on the moral balance in favor of stopping. But we should not build this complication into our account of the difference between positive and negative agency. For the force of this extra factor seems independent of facts about agency. It does not seem to derive from any supposition that, if you stay with the passengers, you will really be taking the train forward or will somehow be party to the fatal action of the train itself.

inaction here than in Rescue III. And this is because of your intentions. In Rescue III you intend an action of the train that in fact causes the man's death, its passing over the spot where he is trapped.²⁵ Not, of course, because he is trapped there. But because the train must pass that spot if the five are to be saved. In Rescue IV, however, things are different. In that case you intend no action of the train that leads to the man's death. The purposes for which you act would be just as well served if the train's brakes were accidentally to apply themselves.

In Rescue III, but not in Rescue IV, the train kills the man *because* of your intention that it continue forward. This implicates you, I believe, in the fatal action of the train itself. If you had no control, but merely wished that the rescue would continue—or if, as in Rescue IV, you had control but no such wish—you would not be party to the action of the train. But the combination of control and intention in Rescue III makes for a certain kind of complicity. Your choice to let the train continue forward is strategic and deliberate. Since you clearly *would* have it continue for the sake of the five, there is a sense in which, by deliberately not stopping it, you *do* have it continue. For these reasons your agency counts as positive.

The surprise in this is that we must bring the distinction between what is intended and merely foreseen into the DDA. But the two doctrines do not therefore merge. As I shall try to show in another paper, the DDE depends on something different—on whether or not a victim is *himself* an intentional object, someone whose manipulation or elimination will be useful. But the victim is not in that way involved in the special kind of positive agency we find in Rescue III. What is intended there is not something for him—that he be affected in a certain way—but some action of an object that (foreseeably but quite unintentionally) leads to his death.

²⁵In Rescue III you intend an action of the train that immediately kills the man. But it would make no difference if, in a variant of the case, you did not intend that the train pass over the spot where the man was trapped, but merely intended that it pass over some nearer part of the track (where that would foreseeably lead to its passing over the fatal spot). Nor would it matter, in a further variant of the case, if the intended action of the train would lead to the man's being killed by some immediate cause other than the train. All that is essential is that you intend some action of the train that you can foresee will cause the man's death.

To the idea of positive agency by action, we must therefore add positive agency by this special kind of inaction. But this is, I think, the only complication we need to build into the doctrine itself. (Other more minor qualifications will be discussed in the next section.) We may now construct the doctrine in stages, starting with some definitions. An agent's *most direct contribution* to a harmful upshot of his agency is the contribution that most directly explains the harm. And one contribution explains harm more directly than another if the explanatory value of the second is exhausted in the way it explains the first.

In the absence of special circumstances involving the actions of objects, an agent's contributions to various effects in the world are those of his voluntary actions and inactions that help produce the effects. So in ordinary cases, his most direct contribution to any effect is the action or inaction that most directly explains the effect. In Rescue I, for example, our most direct contribution to the death of the one is our failure to save him. Our going off to save the five contributes less directly. For it explains the death precisely by explaining the failure to save.²⁶ In Rescue II, on the other hand, our most direct contribution to the death of the man trapped on the road is our act of running him over.

In special circumstances, that is, where harm comes from an active object or force, an agent may by inaction contribute the harmful action of the object itself. This, as we have seen, happens just in case the object harms because the agent deliberately fails to control it and he fails to control it because he wants some action of the object that in fact leads to the harm. Having defined this much, the rest is straightforward. Harmful positive agency is that in which an agent's most direct contribution to the harm is an action, whether his own or that of some object. Harmful negative agency

²⁶We fail to rescue the one *because* we rescue the five instead. But notice that this account implies, in the previously mentioned puzzle cases of boycotting and snubbing (cases where we are unsure whether there is a genuine action by way of an inaction or merely a deliberate inaction), that the agent's most direct contribution to the upshot is an inaction. Grape sales decline because we don't buy grapes, and we don't buy them *because* we are boycotting. Happily, this means that we do not have to decide whether boycotting is a genuine action in order to determine the boycotter's agency in the intended upshot. It will turn out on either hypothesis to be negative.

is that in which the most direct contribution is an inaction, a failure to prevent the harm.

III.

We should now look briefly at certain kinds of cases in which common-sense morality seems to qualify the doctrine as I have just described it, permitting us to harm or even kill someone in order to help others. I am not thinking here of the avoidance of great catastrophes. The doctrine, as already indicated, need not be absolutist. And even in its nonabsolutist form, it cannot contain everything of moral relevance. Special rights to do that which produces harm and special duties to prevent harm must also be factored in. In this way the doctrine has the force of one important *prima facie* principle among others. Rights of competition, to give a familiar example, legitimate certain kinds of harmful positive agency—such as the shrewd but honest competition in which you take away another person's customers. The right to punish is another familiar example. On the other side, special duties to aid may arise from jobs, contracts, natural relations, or from the fact that someone's present predicament was of your making.²⁷ These special duties explain why some instances of negative agency seem no easier to justify than active harmings.

These familiar rights and duties do not require that the doctrine be qualified. They merely oppose it in particular cases. But other situations seem either to require special amendments to my definitions of positive and negative agency or to show that in certain situations the doctrine lacks its usual *prima facie* force. Qualifications of the first sort sometimes seem required where harm arises from the active *withdrawal* of aid. In one kind of case you actively abort a project of rescuing or helping that, knowing what you now know, it would have been wrong to undertake. For example, you stop the train in Rescue III, and the five therefore die. In another kind of case, you remove something from where it would help fewer to where it would help more, for example, a raft that is pres-

²⁷If you have advertently or inadvertently poisoned someone who can yet be saved by an antidote that you actually have, then you seem to be in no moral position to go to the rescue of five others rather than staying to save him.

ently within the reach of one drowning victim but that could be moved to the vicinity of several other victims.²⁸ The object might be your body. You might, for example, cushion the fall of one baby if you stay where you are, but cushion the fall of several others if you move. In all these cases harm comes to someone because you decide to act rather than to do nothing. But because your action is a certain kind of withdrawing of aid, it naturally enough seems to count as negative agency.

In other cases, harmful positive agency seems to lack some of the *prima facie* opprobrium that usually attaches to it. Sometimes this is because the harm would have been avoided but for some blameable fault of the person harmed. Suppose, for example, that the person in Rescue II who blocks the road had been repeatedly warned not to stray where he might interfere with important rescue efforts. If so, we might feel somewhat more justified in proceeding with the rescue (although never, I think, as justified as we feel in Rescue I). People must, after all, accept some responsibility for the predicaments they stupidly and wrongly bring upon themselves.²⁹

In a quite different kind of case, someone may have a special liability to be harmed by a physical or psychological interaction that is generally innocuous and, therefore, of no general moral significance. He might have a rare disease that makes any kind of physical contact very harmful to him. Or he might become dangerously hysterical if we yell in his presence. In such cases we might feel that we could try to save other people from some se-

²⁸It seems important in this kind of case that those who are saved by your action have just as much right to the raft as the one who suffers. Removing it from the reach of its owner would, for example, be very questionable. It also seems important that the person from whom the raft is taken is not already using the raft to save himself. It is one thing to remove it from his reach and quite another to push him off it.

²⁹The responsibility of others also comes in when we know that an action will occasion aggression by a third party—for example, if I know that Jones will murder you if I rescue five of his enemies who are drowning. If it seems that I may proceed with the rescue in this case it is because we shall, quite sensibly, attribute the blame for your death to Jones and not to me. In this kind of situation it is important that the action I undertake is morally pressing. Had Jones threatened to murder you in case I mowed my lawn, my ignoring the threat might well seem a kind of active provocation.

rious danger even if it would mean brushing up against him or yelling. For, unlike standard instances of harmful positive agency, the attempt would not seem to count as an aggression against the victim, since he does not suffer because of any *general or typical* liability to harm. And this seems sensible. Morality must to some degree reflect the standard human condition. In particular, it must be capable of defining a class of presumptively innocent actions.³⁰

Another qualification concerns large public and private projects, like the building of skyscrapers, highways, and dams. We are clearly permitted to help initiate such projects even though we know that in their course some deaths or injuries are practically inevitable. For one thing, the harm is usually remote from what we do. And, more important, the actual harm will generally have been preventable, and its occurrence will be much more directly traceable to the wrongful agency of persons more immediately concerned. It is of course essential that we do not in any way intend the harm that may occur, and take reasonable precautions to prevent it.

In the celebrated Trolley Problem, we seem to find yet another exception to the doctrine's strictures against harmful positive agency.³¹ In this case a runaway trolley threatens five who are trapped on the track where it is now moving. If the driver does nothing the five will die. But he can switch to a side-track where only one person is trapped. Most people think the driver may switch tracks. But switching is positive agency while doing nothing

³⁰But this qualification does not apply to special liabilities created by *external* features of a situation. If driving by Smith's house would set off an explosive device that would blow him up, then driving by, even when it would be necessary to rescue five others, would count as an aggression rather than a failure to help. And this also makes sense. For there seems to be no way in which we can define the class of presumptively innocent actions by prescindng from unusual external circumstances. Removing a ladder is not presumptively innocent when someone is high up on it. And driving down a public road is not presumptively innocent when someone is trapped on it. But entering someone's field of vision (where that sets off no devices, etc.) seems quite different, even where the person will, because of a rare mental illness, be harmed by it.

³¹I believe the case was introduced by Foot in "The Problem of Abortion and the Doctrine of the Double Effect," p. 270. See also Judith Jarvis Thomson, "The Trolley Problem," in *Rights, Restitution, and Risk: Essays in*

appears to be negative agency. So the case looks like a counterexample.

But if we look again, we can see that the driver's passive option, letting the train continue on the main track, is really a form of positive agency. This is because the only possibly acceptable reasons for him not to switch would be to prevent the death of the man on the side-track or to keep clean hands. But the clean-hands motive begs the question; it presupposes that the doctrine does not also speak against not switching. So in deciding the status of his possible inaction we must put this motive aside. This leaves the aim of preventing the death of the man on the side-track. But if the driver fails to switch for this reason, it is because he intends that the train continue in a way that will save the man. But then he intends that the train continue forward past the switch, and this leads to the death of the five. So, by my earlier definitions, his choice is really between two different positive options—one passive and one active.³² And that is why he may pick the alternative that does less harm. Properly understood, Trolley Cases are no exception to the doctrine.

IV.

Perhaps we have found the basic form of the doctrine and the natural qualifications that, when combined with other plausible moral principles, accurately map our moral intuitions. But someone will surely object that intuitiveness and correctness are different things and that intuitions about particular kinds of cases may reflect nothing more than conditioning or prejudice. What we need, therefore, is a more philosophical defense of the doctrine, a rationale that can be called upon to support the intuitions.

Moral Theory (Cambridge, Mass.: Harvard University Press, 1986), pp. 94–116. And Jonathan Glover discusses a fascinating real-life trolley case in *Causing Death and Saving Lives* (Harmondsworth, England: Penguin Books, 1977), pp. 102–103. During World War II British intelligence apparently had the power to deceive the German command about the accuracy of rocket attacks on London. Had they chosen to do so—and they did not—they could have redirected the rockets to less densely populated areas outside the city.

³²This solution to the trolley problem works equally well for versions in which the choice belongs to someone who happens upon a trackside switch.

Foot locates a kind of rationale in the distinction, borrowed from the law but applied to morality, between negative and positive rights. Negative rights are claim rights against harmful intervention, interference, assault, aggression, etc. and might therefore naturally seem to proscribe harmful positive agency, whether by action of the agent himself or by action of some object to which, by strategic inaction, he lends a hand. Positive rights, on the other hand, are claim rights to aid or support, and would therefore seem to proscribe harmful negative agency. Foot's idea seems to be that general negative rights are, *ceteris paribus*, harder to override than general positive rights.³³ And while this seems intuitively correct, it is not obvious why it should be so.

The thesis that negative rights are harder to override immediately implies that negative rights take precedence over positive rights. And it is the thesis of precedence that matters most to us, since it applies directly to circumstances, such as the ones we have been considering, in which the two kinds of rights compete with each other—situations in which the positive rights of one person or group can be honored just in case the negative rights of another person or group are infringed. In Rescue II, for example, the positive rights of the five to be saved from death compete in this way with the negative right of the trapped person not to be killed.

The weakest thesis of precedence would hold that in such oppositions the negative rights prevail just in case the goods they protect (the goods that would be lost if they were overridden) are at least as great as the goods protected by the positive rights (the goods that would be lost if they were overridden). The goods in question are life, health, freedom from injury, pleasure *de facto* liberty, etc.—goods that do not include or presuppose the moral good of respect for any of the rights in conflict.³⁴ All other things being equal, the weakest thesis of precedence would forbid us to

³³See "The Problem of Abortion and the Doctrine of the Double Effect," p. 27. Foot does not actually speak of "general" positive and negative rights. But I think that is what she means. For natural or contractually acquired "special" positive rights may sometimes bind as strongly as general negative rights. We saw, for example, that a private lifeguard in Rescue I might not be permitted to leave to save the five.

³⁴In presenting versions of the precedence thesis, I am supposing (*contra* John Taurek in "Should the Numbers Count?" *Philosophy and Public Affairs* 6 (1977)) that the numbers do count—for example, that saving two

kill one person to save another, but would permit us to kill one in order to save two.

A very strong thesis of precedence, on the other hand, would rule out any infringement of certain very important negative rights (for example the right not to be killed or the right not to be tortured) no matter what positive rights were in competition with them. This would still allow positive rights protecting more important goods to prevail over negative rights protecting less important goods—would permit us, for example, to knock one person down in order to save another from serious injury. But it would not permit us, for example, to kill or torture one to save any number of others even from death or torture.

A perhaps more plausible intermediate thesis would hold that no negative rights are absolute, but would accord to the most important ones considerably more force than they have on the weakest thesis. Such a view might well accommodate the ordinary thought that while someone may not be killed to save five, he might be killed to stave off the kinds of disasters that consequentialists dream up. It might go on to state some kind of criterion for when negative rights must give way; or it might, in Aristotelian fashion, leave the matter to moral perception.³⁵

If, on the other hand, negative rights do not take precedence over positive rights then either the reverse is true or neither takes precedence over the other. If positive rights actually take precedence, then we might, as seems absurd, kill two to save one. Suppose one person is drowning and two are trapped on the road. A morality that permitted us to run over and kill the two in order to save the one seems not only odious but incoherent. For once we have decided to kill the two, we have placed them in at least as much danger as the one was in originally. And that would presumably activate their positive rights to be saved from their predica-

lives generally does twice as much good as saving one. I am also supposing that goods of different kinds (for example, preservation of life and relief from suffering) can be compared and at least roughly summed up, and that in cases of conflicting rights we can make at least a rough comparison of the overall good protected by the rights on each side of the conflict.

³⁵Or it might include a criterion that itself requires intuition to apply—by claiming, for example, that a negative right may be justifiably infringed just in case it would be contemptible of its possessor to insist on it. That is the kind of criterion that I find attractive.

ment—rights that would collectively outweigh the positive rights of the one who is drowning.

If there is going to be precedence, it clearly has to be precedence of negative rights. But this leaves open the possibility that neither kind of right takes precedence over the other, that is, that in the competitions we are considering the rights protecting the greater balance of good should, *ceteris paribus*, prevail. In such a moral system the person trapped on the road in Rescue II could not with moral authority object to our running over and killing him. For we shall be saving five others each of whom values his life just as much as he values his. This moral system is perfectly coherent. But it has unappealing aspects.

In such a morality the person trapped on the road has a moral say about whether his body may be destroyed only if what he stands to lose is greater than what others stand to gain. But then surely he has no real say at all. For, in cases where his loss would be greater than the gain to others, the fact that he could not be killed would be sufficiently explained not by his authority in the matter but simply by the balance of overall costs. And if this is how it is in general—if we may rightly injure or kill him whenever others stand to gain more than he stands to lose—then surely his body (one might say his person) is not in any interesting moral sense *his*. It seems rather to belong to the human community, to be dealt with according to its best overall interests.

If it is morally his, then we go wrong if, against his will, we destroy or injure it simply on the ground that his loss will be less than the gains of others. The same is true of his mind. If we may rightly lobotomize or brainwash him whenever others will gain more than he will lose, then his mind seems to belong not to him but to the community. There is an obvious parallel here with his different, and much less important, relation to his property. An object does not belong to him if he may have and use it, and others may not take it from him, only as long as his keeping it would be better for him than his losing it would be for them.³⁶ Whether we are speaking of ownership or more fundamental forms of possession, something is, morally speaking, his only if his say over what may be

³⁶And something similar holds for damage. You don't own something if others may damage it whenever that is best for all concerned.

done to it (and thereby to him) can override the greater needs of others.³⁷

A person is constituted by his body and mind. They are parts or aspects of him. For that very reason, it is fitting that he have primary say over what may be done to them—not because such an arrangement best promotes overall human welfare, but because any arrangement that denied him that say would be a grave indignity. In giving him this authority, morality recognizes his existence as an individual with ends of his own—an independent *being*.³⁸ Since that is what he is, he deserves this recognition. Were morality to withhold it, were it to allow us to kill or injure him whenever that would be collectively best, it would picture him not as a being in his own right but as a cell in the collective whole.³⁹

This last point can be illustrated not by thinking of bodies or minds but of lives. The moral sense in which your mind or body is yours seems to be the same as that in which your life is yours. And if your life is yours then there must be decisions concerning it that are yours to make—decisions protected by negative rights. One such matter is the choice of work or vocation. We think there is something morally amiss when people are forced to be farmers or

³⁷Reference to the specific moral relation that I have in mind (in saying that someone's body and mind are his and not the community's) is made most naturally by a particular moral use of the possessive pronoun. This makes repeated reference awkward, and tempts me to talk in ways that are potentially misleading. I have spoken of a person's mind as belonging to him and have drawn an analogy with property. But both moves are dangerous. The intended sense of "belong" derives from the special use of the possessive. And the analogy with property is, as indicated, inexact. Both relations ground rights of say in what is to be done, but a person's mind or body are definitely not property—not even his property.

³⁸I mean here to invoke the ordinary sense of "being," in which human persons, gods, angels, and probably the higher animals—but not plants, cells, rocks, computers, etc.—count as beings.

³⁹It would make no difference, I think, if the overall good of the whole were thought to be a mere sum of the good of its parts—that is, if the whole were regarded as a mere colony without a morally significant higher-order function of its own. To deny the precedence of negative rights would still be to limit a person's moral protections precisely by this test: whether or not granting the protections would best serve the collective good. It would be to suppose that he may rightly be killed or injured if the cost to him does not outweigh the sum of the benefits to others. And this seems to me a clear enough way in which he would be regarded, morally, as a cell in the collective whole.

flute players just because the balance of social needs tips in that direction. Barring great emergencies, we think people's lives must be theirs to lead. Not because that makes things go best in some independent sense but because the alternative seems to obliterate them as individuals. This obliteration, and not social inefficiency, is one of the things that strikes us as appalling in totalitarian social projects—for example, in the Great Cultural Revolution.

None of this, of course, denies the legitimate force of positive rights. They too are essential to the status we want as persons who matter, and they must be satisfied when it is morally possible to do so. But negative rights, for the reasons I have been giving, define the terms of moral possibility. Their precedence is essential to the moral fact of our lives, minds, and bodies really being ours.

But it might be objected that the weakest thesis of precedence would give us some degree of moral independence, and at the same time would let us do the maximum good, honoring as many positive rights as possible. On that thesis, it would not be proper to kill one person to save another who is equally happy and useful—it would not be proper, say, to flip a coin. But it could be right to kill one to save two or even five to save six. Why then adopt a stronger thesis? The answer, I think, depends on how important the relevant forms of legitimate control are to us—the extent to which we wish to belong, in the sense under discussion, to ourselves.⁴⁰ And this might depend on the aspect of ourselves in question.

We feel, I believe, most strongly about assaults on our minds. Here most of us are far from minimalists about the precedence of negative rights. The idea that against our will we could justifiably be brainwashed or lobotomized in order to help others cuts deeply against our sense of who and what we are. Here it seems the sense of our own rightful say leads almost to absolutism. We feel less strongly about our persons (at least those parts that do not directly affect our minds) and labor. But even here we wish, I think, to have a kind of defensive say that goes far beyond the weakest

⁴⁰I am not claiming that any person or persons have actually designed morality with an eye to giving themselves the degree of say they find fitting. But I do think that light can be shed on the (timeless) content of morality by considering the importance to us of what would be realized or unrealized in the design of various moral systems.

thesis of precedence. A system that gave you some authority over what might be done to you but allowed us to kill or injure you whenever that would even slightly maximize the overall good would seem a form of tokenism.

It must be said that something like the precedence of negative rights can be accepted by a certain kind of consequentialist—one who thinks that a person's having an effective say over what is done to him (but not over what is done to others) is, in itself, a kind of good that can be added to the more familiar goods of life or happiness.⁴¹ This kind of consequentialism would grant each of us a kind of special authority against interference. But it is unclear that it would thereby give us the moral image of ourselves we think fitting. For it locates the ultimate ground of proper deference to a person's will in the fact that such deference maximizes the general balance of good. In such a system, it is not so much his right to have his way that really matters as the general goodness of letting him have his way.

A consequentialist might reply that anything other than a consequentially grounded system of rights leads to absurdities, and that in praising the virtues of a rights-based morality I can be saying no more than that there is value in the social influence of such a system—that it is good if people's rights are respected and bad if they are violated. But circumstances can arise in which respecting someone's negative rights will lead to an abuse of the negative rights of others. And in at least this kind of case it would be incoherent, the consequentialist will insist, to suppose that negative rights can override their positive counterparts.⁴² Suppose B and C will be murdered unless we murder A. A has a negative right against our murdering him, and B and C have positive rights that we help prevent their being murdered. If the ground of the system of rights lies in the value of respect for (or at least nonvio-

⁴¹Amartya Sen makes room for what he calls goal rights in "Rights and Agency," *Philosophy and Public Affairs* 11 (1982), pp. 3–39.

⁴²Samuel Scheffler develops such an argument in *The Rejection of Consequentialism* (Oxford, England: Clarendon Press of Oxford University Press, 1982), pp. 80–114. Sen, in "Rights and Agency," Section VI and VII, tries to make room within a consequentialist framework for kinds of agent-relativity that would undermine the argument. But I find these agent-relative features poorly motivated as elements of a possible consequentialism.

lation of) rights, then surely the positive rights of B and C must prevail. For only by murdering A can we maximize the value that the entire system aims at.

But this objection misses the mark. The value that lies at the heart of my argument—the appropriateness of morality’s recognizing us as independent beings—is in the first instance a virtue of the moral design *itself*. The fittingness of this recognition is not a goal of action, and therefore not something that we could be tempted to serve by violating or infringing anybody’s rights. It is also true, of course, that we think it good if people actually respect each other’s rights. But this value depends on the goodness of the moral design that assigns these rights. It is not that we think it fitting to ascribe rights because we think it a good thing that rights be respected. Rather we think respect for rights a good thing precisely because we think people actually have them—and, if my account is correct, that they have them because it is fitting that they should. So there is no way in which the basic rationale of a system of rights rules it out that a person might have a right not to be harmed even when harming him would prevent many others from being harmed in similar ways.

The rationale that I have proposed is anticonsequentialist not only in its assignment of priority to negative rights, but also, and more fundamentally, in its conception of the basic social function of morality. For consequentialism, it seems fair to say, the chief point of morality is to make things go better overall—to increase average or total welfare within the human community. But on the view presented here, an equally basic and urgent moral task is to define our proper powers and immunities with respect to one another, to specify the mutual authority and respect that are the basic terms of voluntary human association. The doctrine we have been discussing addresses this task directly. And this is why it is far more than a casuistical curiosity. Whether we ultimately agree with it or not, we should recognize that, in giving each person substantial authority over what can rightly be done to him, the doctrine conveys an important and attractive idea of what it is to be a citizen rather than a subject in the moral world.

University of California, Los Angeles