

INTENTION AND PERMISSIBILITY¹

by T. M. Scanlon and Jonathan Dancy

I—T. M. Scanlon

ABSTRACT It is clearly impermissible to kill one person (or refrain from giving him treatment that he needs in order to survive) because his organs can be used to save five others who are in need of transplants. It has seemed to many that the explanation for this lies in the fact that in such cases we would be intending the death of the person whom we killed, or failed to save. What makes these actions impermissible, however, is not the agent's intention but rather the fact that the benefit envisaged does not justify an exception to the prohibition against killing or the requirement to give aid. The difference between this explanation and one appealing to intention is easily overlooked if one fails to distinguish between the prospective use of a moral principle to guide action and its retrospective use to appraise the way an agent governed him or herself. Even if this explanation is accepted, however, it remains an open question whether and how an agent's intention may be relevant to the permissibility of actions in other cases.

I

Does the permissibility of an action depend on the intentions of the agent? In particular, can the fact that the agent intended a certain outcome either as an end or as a chosen means to that end make an action that would otherwise have been morally permissible morally impermissible? Defenders of the Doctrine of Double Effect believe that the answer is yes, and a number of examples seem to suggest that this answer is correct. The best known of these are a series of three hospital cases.

Drug Shortage. There are five people in Room B, and one person in Room A, all of whom have the same disease, and all of whom will die if not treated soon. There is enough medicine on hand to cure all five of the people in Room B, but since the

1. This paper developed out of material presented in my moral philosophy seminar at Harvard University in the Spring Term of 1999, and it owes a great deal to the helpful comments and suggestions of the members of that group. Versions of the paper was presented at the University of York and at the New York University Colloquium in Law, Philosophy and Political Theory. I am grateful to the members of both audiences for their comments and criticisms, but I have not been able to respond to all of these in the present version. Kent Greenawalt provided valuable suggestions at an earlier stage.

person in Room A has a more advanced case, it would take all of the available supply just to save him. Everyone agrees that in this case it is at least permissible to use all of the medicine to save the people in Room B.

Transplant. The five people in Room B are in need of organs—one needs a heart, two need a lung and two a kidney—and they will all die if they are not given transplants soon. Unfortunately, no organs are presently available. But there is a person in Room A, in for a check up, who could be given a lethal injection instead of the inoculation he is expecting, thereby making his organs available to save the five. All agree that this is impermissible. This might be explained by saying that although it is permissible to let one die in order to save five, as in *Drug Shortage*, what is proposed in this case is to kill the one in order to save the five, and this is not permissible.

Drug/Transplant. The people in Room B are the same as in *Transplant*, and the person in Room A is the same as in *Drug Shortage*. If this person dies of his illness, his organs will not be damaged and thus can be used to save the five. Is it permissible to refrain from giving him the available medicine in order to save the greater number? Everyone with whom I have discussed these cases agrees that it is not. Why not?

One answer would be that, as in *Transplant*, this would be a case of killing one to save five. But this cannot be the basic explanation. What it is proposed to do to the one in this case is the same as in *Drug Shortage*—to refrain from giving him the available drug—and it was said in that case that this was not killing. If it is killing in this case, then this is because ‘killing’ is being used as a moral notion—that is to say, it is killing because it is wrong (for some other reason) rather than being wrong because it is killing. What, then, is this other reason? One explanation that naturally comes to mind is that whereas in *Drug Shortage* the death of the one is merely a foreseeable consequence of giving the drug only to the five, in *Drug/Transplant* the death of the one is intended: what is proposed is to withhold the drug from the one precisely in order to get him dead right away, so that his organs will be available for transplant in time to save the five. This is, on the face of it, extremely plausible as an explanation of the difference between the two cases, as a result of which what is proposed in *Drug Shortage* is permissible, but what is proposed in *Drug/Transplant* absolutely impermissible.

One might suspect that our reactions to these cases depend on ideas that are peculiar to the hospital setting—on ideas about the duties that hospital personnel owe to patients. I will return to this possibility below, but it is worth noting here that similar problems can arise in cases that do not involve hospitals or the obligations of doctors.

Rescue I. As I am driving home, I hear on my citizens band radio that a car is stalled along a seldom-traveled road that I could easily take. The car is delivering medicine to someone who will die unless he receives it within the next few hours. I could easily take that road and restart the stalled car. Clearly I should do so.

Rescue II. Same as the previous case except I also hear that along another road I could take there is a stalled car that was taking medicine to five people in equally urgent need. There is not enough time for me to go to the aid of both cars. Clearly it is at least permissible for me to aid the latter car, so as to save five rather than only one.

Rescue/Transplant. Same as *Rescue I* except that I know that there are five people in urgent need of transplants who will be saved if the patient awaiting the medicine dies very soon, as he will if I do not go to the aid of the stalled car. May I therefore refrain from aiding the car? It seems that I cannot. Why not? As before, one appealing explanation is that in this case, but not *Rescue II*, I would be deliberately allowing the one to die, as a means to saving the five.

At least I have found this explanation immediately appealing. But there are well-known problems with it. First, no one has, to my knowledge, come up with a satisfying theoretical explanation of why the fact of intention in the sense that is involved here—the difference between consequences that are intended and those that are merely foreseen—should make a moral difference. Second, there are cases in which applying this distinction seems to give the wrong answer. For example, in the standard trolley problem case, it seems permissible to switch the runaway trolley onto a track on which it will hit only one person rather than allow it to continue straight ahead and hit five. But it also seems permissible to turn the trolley in the Loop case, proposed by Judith Thomson, in which if the trolley does not strike the one it will continue around the loop and hit the five from the other

side.² These cases seem to differ in just the way I have described: in the second, but not the first, one switches the trolley only because it will, by hitting the one person, be prevented from hitting the five. So it seems that the one's being hit by the trolley is intended as a means to the end of saving the five.³ Perhaps the distinction between intended and merely foreseen consequences, even though it makes a moral difference in the cases described above, does not make a difference in these cases. But this needs to be explained.

Thomson also presents, as counter-examples to the kind of analysis we are considering, cases involving the use of lethal drugs for pain relief. Suppose that a patient is fatally ill and in great pain. The only course of medication that will relieve this pain will also cause the patient's death. Suppose that the patient wants to take this drug. Does the permissibility of administering it depend on the doctor's intention in doing so—specifically on whether the doctor intends to relieve the pain by causing the patient to die or intends to relieve the pain by giving the drug which will, inevitably, also cause the patient's death? Thomson says, plausibly, that this would be absurd.

The idea that this would be absurd may gain support from the thought that it is not a bad thing, morally speaking, for a person in such circumstances to die sooner rather than later. But Thomson's objection retains its intuitive force in other cases in which this is not an issue. It is plausible to claim, for example, that it can be permissible in wartime to bomb a munitions factory even though this is certain to kill some civilians living nearby, but that it would not be permissible to kill the same number of civilians just as a way of undermining public support for the war even if this was just as likely as the destruction of the munitions plant to hasten the end of the conflict. But even if these two strategies differ in moral permissibility, it does not seem plausible to say that the moral difference between them is a matter of the outcomes intended by the agents involved. If the commander were to come to you, as Prime Minister, describe a planned raid that

2. See Thomson, 'The Trolley Problem', in her collection, *Rights, Restitution and Risk*, edited by William Parent (Cambridge, MA: Harvard University Press, 1986), pp. 101–102.

3. In 'The Law of Triple Effect', Frances Kamm questions whether it follows from the fact that one does X only because it will lead to Y that one intends Y.

would be expected to destroy a plant and kill a certain number of civilians, thereby probably undermining public support for the war, and ask whether you thought this was morally permissible, you would not say, 'Well, that depends on what your intentions would be in carrying it out. Would you be intending to kill the civilians, or would their deaths be merely an unintended but foreseeable (albeit beneficial) side effect of the destruction of the plant?'⁴

If the idea that the permissibility of an action depends on the agent's intention in performing it is as implausible as this makes it sound, why should it have seemed so appealing in the cases with which I began? Thomson suggests that the persistence of this idea depends on 'a failure to take seriously enough the fact—I think it is plainly a fact—that the question whether it is morally permissible for a person to do a thing is just not the same as the question whether the person who does it is thereby shown to be a bad person'.⁵ As she says, a doctor who dislikes her patient and administers a lethal dose of pain killer, relishing the thought that this will be the last of him, is acting from a bad motive. This is a serious moral fault. But it does not follow that it is impermissible for her to administer the drug (or that the patient should have to wait until a different doctor, with better intentions, comes on duty.)

There is something right about this suggestion. What an agent takes as counting in favour of a course of action does seem to bear on a moral assessment of that person in a way that it may not bear on the permissibility of what the agent does. There is such a thing as doing the right (or a permissible) thing for reasons that show one to be a bad person. But it is not clear that a failure to take this distinction seriously can explain the appeal of the idea we have been considering. As Jonathan Bennett has pointed out, although what a person takes as ultimate ends and what he or she is willing to countenance in pursuit of those ends are both clearly relevant to a moral assessment of that person, it is not clear that the distinction between intended and merely foreseen consequences—that is, between what a person adopts as means to

4. This example is derived from one suggested to me by Thomson.

5. 'Physician-Assisted Suicide: Two Moral Arguments', *Ethics* 109 (1999), p. 517.

his or her ends and what a person is merely willing to countenance for the sake of achieving those ends—makes a difference to such an assessment.⁶ Perhaps it does make a difference, but this is not evident, and it has to be said that in the case Thomson cites our negative assessment of the doctor who out of malice or desire for revenge seeks the death of her patient seems to depend on the doctor's ultimate end rather than merely her chosen means.

My aim in this paper is to propose a different way of making sense of these phenomena. In particular, I will offer an account of what it is that makes actions like those involved in *Transplant* and *Drug/Transplant* impermissible, and point out how it is tempting to confuse this feature with a fact about the agents' intentions in these cases. I agree with Thomson that this confusion involves a failure to distinguish clearly between features that make an action wrong and descriptions of the flaw that an agent exhibits in performing it. But I will offer a more detailed and slightly different explanation of this mistake.

II

As background for the account I am going to offer, I will begin with some general points about the relations between intention and moral assessment. 'Intention' is commonly used in wider and narrower senses. When we say that a person did something intentionally, one thing we may mean is simply that it was something that he or she was aware of doing, or realized was a consequence of his or her action. This is the sense of 'intentionally' which is opposed to 'unintentionally'. To say that you did something unintentionally is to claim that it was something you did not realize you were doing. But we also use 'intention' in a narrower sense. To ask a person what her intention was in doing a certain thing is to ask her why she was doing it—which of the various things she realized she was doing were things she took to count in favour of so acting. This narrower sense is much closer to (perhaps even the same as) the idea that is involved in the distinction between the consequences of one's action that are intended (as ends or chosen means) and those that are merely foreseen.

6. Jonathan Bennett, 'Morality and Consequences', *The Tanner Lectures on Human Values*, Vol. II, S. McMurrin, ed. (Salt Lake City: University of Utah Press, 1981), p. 99.

Both the wider and the narrower senses of ‘intention’ have something to do with an agent’s view of the reasons bearing on his or her action. This is most obvious in the case of the narrower notion: one’s intention, in this sense, is an aspect of one’s action that is crucial to the reason one takes oneself to have to do it. But it is also true that if one does something intentionally in the broader sense of that term—if one is aware that in acting one is doing it—then even if one does not take this aspect of what one is doing to provide a reason for so acting, one at least does not (insofar as one is not acting irrationally) take it to constitute a conclusive reason against acting in that way. In either case, then, to say what an agent does intentionally, or intends, is to say something about his or her assessment of the reasons that bear on acting in that way. It is this connection with what an agent took to be reasons for (or against) acting a certain way that makes these notions of intention relevant to the moral appraisal of agents.

Moral criticism is appraisal of the way we govern ourselves—of what we do or do not take into account in regulating our conduct and of the way we take it into account as reasons for and against acting in certain ways. Moral criticism is generally at least hypothetically retrospective. That is to say, it is predicated on assumptions (perhaps hypothetical) about how a person governed him or herself—what he or she took into account and took as a reasons for and against acting as he or she did. Moral principles, on the other hand, although they can be employed retrospectively in criticism of this kind, are in the first instance prospective—they specify the ways in which agents are to go about deciding what to do.

People often speak of moral principles as if they were specific directives, telling us what to do and what not to do, but this description is quite misleading: principles rarely, if ever have this form.⁷ Some principles do state what are normally conclusive reasons for certain courses of action and against others. For example, we should, in most circumstances, take the fact that we have promised to do something as a conclusive reason for doing it, and the fact that a certain course of action can be seen to set

7. The following discussion of principles and of our understanding of the exceptions they incorporate draws on a fuller account in my book, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1998), pp. 197–202.

in motion a chain of causes leading directly to someone's death normally counts as a conclusive reason against it. But even these familiar and important moral requirements have exceptions. Breaking one's promise is sometimes the right thing to do; there are circumstances (such as self-defense) in which killing is justified; and in other circumstances (such as that described in the Trolley Problem) it is permissible to do what one has every reason to believe will involve killing an innocent person.

Although it is sometimes unclear whether a certain consideration does or does not count as justifying an exception to some principle, there are also many cases in which it is quite clear to anyone who understands a principle whether a consideration justifies an exception or not. But it is generally very difficult—I would say often impossible—to spell out a principle in a way that would enable one to decide, without appeal to judgment, even these intuitively clear cases. Any formula we could articulate would have 'other things equal' clauses, or qualifiers such as 'sufficiently important', which call for substantive moral judgment. This raises the question of what we are doing—what kind of thinking we are engaging in—when we 'discover' that certain considerations clearly do, or do not, constitute an exception to some moral requirement or prohibition, or when we find that it is unclear whether they do or do not.

My own view is that what we are doing in such cases is asking what form the relevant principles would have to take if they are to be principles that no one could reasonably reject as authoritative guides to conduct. We begin from an (incompletely specified) idea of a general requirement or prohibition that we believe meets this test. We then use our understanding of the reasoning behind this principle to specify it further, determining whether the consideration in question would or would not need to be recognized as marking out an exception to this general prohibition or requirement in order for the principle to be one that no one could reasonably reject.⁸ So, for example, in the case of

8. In emphasizing the specification of principles, my view is close the one put forward by Henry Richardson in 'Specifying Norms as a Way to Resolve Concrete Ethical Problems', *Philosophy & Public Affairs*, 19 (1990), pp. 279–310. In Richardson's version, however, the need for specification arises from conflict between principles and involves modifying them. On my account, only one principle need be involved and the process is typically one of figuring out what the principle requires rather than modifying it.

promises, we begin with the idea that principles that no one could reasonably reject could not leave people free to create expectations of the kind involved in promises and then disappoint these expectations whenever they wished to. But, on the other hand, a principle that required promises to be kept 'no matter what' would be so constraining as to be reasonably rejectable. In order to decide whether a given consideration—such as a certain cost to the promisor—constitutes good grounds for breaking a promise we thus need to consider whether a principle that did not incorporate it could be rejected as too constraining, and also whether one that did recognize it would seriously undermine the point of promising. (The question of whether a given cost to the promisor counts as an exception to the requirement to keep one's promises will thus depend not only on the magnitude of that cost but also on other factors, such as whether it is something that could have been foreseen at the time the promise was made.) In some cases, the answer will be quite obvious once we become clear about the issues at stake. In others the answer may remain quite unclear, perhaps because the considerations involved are so evenly balanced that it is a 'judgment call' or because the considerations in question are so unfamiliar that it is unclear what the effect would be of recognizing them as an exception.

This account of the process of moral decision in new and possibly difficult cases is to be contrasted with various other models including, for example, the idea that moral thinking employs more fully specifiable principles which always apply but which must be balanced against one another in cases of conflict, the outcome depending on their relative moral 'weight'. An example of this approach would be the idea that we have both a duty to aid and a duty not to harm but that the latter has greater weight, and prevails when duties of the two kinds conflict. In my view, by contrast, the question we should ask is not about the relative weight of conflicting principles but rather about how the relevant principles themselves are best understood. So, for example, in the case just mentioned, the questions would be whether the need to aid someone would be recognized as an exception to the general prohibition against killing, in a proper understanding of that principle, and whether the fact that one could aid one person only by killing another would be recognized as justifying an exception to the duty to aid, when that duty is properly understood. (I am inclined to think, in fact, that clear and helpful

answers are more likely to be found only when the proposed exceptions are defined in more specific terms.)

The account I have proposed involves balancing. This is not, however, balancing of conflicting principles but rather of considerations that figure in the justification of individual principles and determine their content. In addition, the outcome of the process I have described depends not only on the relative weight of these considerations but also on the ways in which recognizing them would interact with other reasons (as the example of promising, considered above, illustrated.)

Much of this section has been devoted to the thesis that plausible moral principles do not merely state general requirements but also incorporate exceptions to these requirements. I have concentrated on this apparent truism for two reasons. The first is to make the point that almost all principles (not just those stating ‘imperfect duties’) are best understood as making (often rather complex) claims about the considerations that we should recognize as reasons for action. The second is to call attention to the importance of exceptions in our moral thinking, and to sketch an account of how we go about interpreting these in difficult cases.

III

With all of this as background, let me turn now to the question of the role of intention, and, more specifically, of the idea of the intended consequences of an action, in moral principles and moral criticism. Since, as I argued above, these notions concern or have implications for what an agent took to be reasons for or against his or her action, they are certainly the kind of things that moral criticism and moral principles, as I have described them, are concerned with. The question is whether they figure in moral argument in a specific way: as considerations that can make the difference between an action’s being permissible and its being impermissible.

A moral principle, I have said, identifies certain considerations that are to count for, or against, acting in a certain way. An action is impermissible if, under the circumstances, there are such reasons that count decisively against it; otherwise, it is permissible. It is required if, under the circumstances, there are reasons

that count decisively in its favour. The relevant principles say what these reasons are and that they are decisive, and if we accept such a principle, and follow it, we will treat the reasons it says are decisive as decisive. If we do not do this, then we act wrongly. But what makes our action wrong is not the reason on which we acted but, rather, the features of the action and its circumstances that the principle identifies as decisive reasons against it.

Consider an example. Suppose that I have promised to do something, and that under the circumstances this counts as a decisive reason for doing it. In particular, the fact that I could benefit financially from breaking the promise is not a sufficient reason to fail to keep it. But suppose I break it anyway, in order to get this benefit. In describing what was defective about my action, you might say that I acted wrongly in taking my own advantage as sufficient reason to break my promise. But, at a more fundamental level, what made my action wrong was not the reason I acted on but the reason counting against so acting. The act was wrong because the fact that I promised made it the case, under the circumstances, that I should do the thing in question. This is even clearer when we view the case prospectively. Suppose I ask, while deciding what to do, 'Must I do what I promised? Why shouldn't I choose the other course, since it would be more profitable?' You would not reply, 'That would be wrong, because you would be aiming at your own benefit (or acting for the sake of your benefit.)' What would make my action wrong is, rather, the fact that I promised and that that, under the circumstances, is a decisive factor.

To take another example, suppose that you are in mortal danger, and I could easily save you, but I would benefit from your death. If I were to fail to save you because, say, I want to have one less rival to contend with, I would be acting wrongly. But, if we consider the matter from my point of view when deciding how to act, what makes this action wrong is not that I would be acting with the wrong intention but the fact that there is a decisive reason to save you, and the fact that the advantage I would gain does not count against this. This is so even though we might also say, truly, when describing the faultiness of my action retrospectively, that what was wrong with my decision was that I took my self-interest as sufficient reason to allow you to die.

This explains what is so odd about the cases that Thomson describes, in which an agent's intention is said to make the difference between permissible and impermissible action. Consider the example of strategic bombing versus terror bombing. The example presupposes that the relevant principles governing the conduct of war are something like this. In time of war one is permitted to use destructive and potentially deadly force of a kind that would normally be prohibited. But this is permitted only when it is justified by a military objective, and only provided that one takes sufficient care to minimize civilian casualties and that the expected casualties are 'proportional' to the importance of the military objective. The example we are considering presupposes that the destruction of the munitions plant is a military objective in the relevant sense. The raid on it is therefore permissible if (and only if) the conditions just listed are fulfilled: if adequate precautions are taken and the expected civilian losses are 'proportional'. The civilians who could be killed in a terror raid, however, are not 'military targets' even if killing them would hasten the end of the war. So terror bombing is not permissible, because the relevant justification is lacking, not because of the intentions of the pilots who carry it out or the commanders who order it. The raid that hits the plant and kills enough people to demoralize the country is permissible just in case it is permissible considered as a raid on the plant alone. It is tempting to say that the impermissibility of terror bombing depends on the intentions of those who plan it or carry it out because it is tempting to move from the true premise that what is morally faulty about their reasoning lies in the fact that they take demoralizing the public as a consideration that justifies the bombing to the false conclusion that this fact about their reasoning is also what makes their action wrong. But it is a mistake to treat the fact that in acting a certain way we would be acting on certain reasons as the consideration that we should take as counting decisively against so acting. This is the absurdity that Thomson identifies.

The rescue/transplant case can also be seen as turning on the existence of general reasons against a course of action and the inadequacy of purported countervailing reasons rather than on the intentions of the agent. The guiding principle in this case says, roughly, that under such conditions we should take the fact that we could easily prevent something very bad from happening

to someone (prevent the invalid from dying by taking the road that will bring us to the stalled delivery van) as a strong reason for doing that. The question is whether this principle must recognize the fact that others would benefit from this person's death as a consideration that justifies an exception to this general requirement.

The principle would take this form, according to contractualism, if those who would benefit from this exception could reasonably reject a principle that did not include it. So the question is: when are the benefits flowing from a person's death something that others have the kind of claim on that is involved in their being able to raise such an objection?

Sometimes people do have such a claim. For example, in self-defence cases killing or disabling you may be the only way for someone to prevent you or your body from killing him. Here the benefit of your injury or death—the non-invasion by you of this person's body—is something that that person can hardly be denied a claim over. Recognizing this claim involves very little diminution of *your* status as a person. But things are quite different if the benefit in question is the use of your organs. To recognize this benefit as a ground for insisting on exceptions to principles protecting our lives would be to take the view that our bodies are resources over which everyone has, at the outset at least, the same claim. It would thus be inconsistent with the idea that each person has a special claim to and authority over his or her own life and body, an idea which I take to be itself one that no one could reasonably reject, and thus to be a notion that shapes other contractualist arguments.

I do not, obviously, have a full defence of this idea. Much more would need to be said to describe the shape of this constraint on the benefits that can figure in defence or rejection of principles. I am inclined to believe, however, that something like this is the best explanation of why the requirement to aid should have the structure I suggest (why the fact that if the person dies others could be saved with his organs should not count as justifying an exception to this requirement.) But my present point is independent of this particular explanation. It is rather that if the requirement to rescue does have the structure I have suggested, then this provides an account of the rescue/transplant case which does not turn on the agent's intention, but which explains why

that intention might seem relevant if one failed to distinguish carefully enough between the fact that an agent's reasoning was faulty and the facts about reasons that make it faulty.

Even if this is the best explanation of these cases, and of the oddness of the appeals to intention as a determinate of permissibility in the cases that Thomson cites, it does not follow that there are no cases in which the fact that the agent has a certain intention (or the fact that the agent takes certain things as counting in favour of his or her action) is the consideration that makes an action wrong. Showing that this is never the case would take a more general argument, and I think it is unclear whether such an argument could succeed. The way to approach this general question within the framework I am employing is to ask what reason either agents or those who would be affected by their actions would have to insist on, or to reject, principles that made the permissibility of an action turn on the agent's intentions.

IV

According to the view I have been advocating, the content of a principle—the set of considerations that it tells agents to count for or against an action—arises out of the interaction between the reasons that those who would be affected by that principle in various ways have for favouring or rejecting it. The question at issue is how the interaction between these reasons might result in a principle that directs us to take the fact that an action would involve acting with a certain intention—for example, that it would involve intending rather than merely foreseeing a certain harm—as counting (perhaps decisively) against that action. This might happen in either of two ways: such facts about intention might figure in the content of a principle because of the importance that agents have reason to attach to actions with a certain intention, or it might arise from the significance that facts about an agent's intentions have for those who are affected by an action. For reasons of space, I will concentrate here on the latter possibility, which seems to me more promising.

There certainly are cases in which we have reasons to care about the intentions of those with whom we interact, and in which the fact that people have certain intentions (in a broad

informal sense of 'intention') can therefore make it wrong for them to act in certain ways. If, for example, I am telling a close friend about my plans to run for office and describing the strategy I intend to use, it would be wrong for her to let this conversation continue if she in fact intends to run for the same post. Similarly, suppose I am distraught about my personal life, and am telling my troubles to what I take to be a sympathetic friend, but that he in fact regards me and my troubles as ludicrous, and is thinking, while he listens, about how he will turn this into a comic novel. Even if nothing comes of this intention, I may have something to complain about in this interaction. If I am not already presuming too much in asking this person to listen to me, it would be wrong of him to continue to do so if he could easily avoid it.

These cases are not, however, ones in which what makes an action wrong, at the most basic level, is the presence of a certain intention. Rather, in both cases I am entering into a certain conversational relationship on the assumption that the other party has certain intentions and attitudes. It is clear between us that I would not be revealing myself in certain ways if I did not believe this assumption to be correct. The wrongness of what the other party does, when it is wrong, lies in continuing to participate in the conversation on this false basis—allowing me to continue to do something that I would not consent to do if I understood what the situation actually was.

One might proceed along similar lines to argue for the importance of intention, in the more specific sense of intended versus merely foreseen consequences, in the hospital cases with which I began. The idea would be that one enters a hospital, putting oneself in the care of its staff and permitting them to manipulate and even invade one's body in ways that are not generally allowed, only because one has certain expectations, which one takes to be shared, about how one will be treated. These include the understanding that hospital employees will use approved medical techniques with the aim of improving one's health. It is understood that the hospital has similar duties to other patients, whose needs may have to be balanced against one's own—one may not, for example, be able to get an operating room right away if it is needed for others. (This limitation explains cases like the drug shortage example.) But it is also understood that hospital

staff will not, without one's permission, aim at making one's medical condition worse, even if other patients would benefit from this. It is this understanding that is violated in the drug/transplant case.

Several comments about this explanation are in order. Just as in the preceding cases, and in contrast to the Principle of Double Effect itself, this analysis does not make intention (or intending certain consequences as ends or means) a basic wrong-making feature. The basic moral idea is the wrongness of violating the understanding (supposing it to be reasonable) on the basis of which patients have placed themselves in the care of hospital employees and given them permission to manipulate and invade their bodies. Intentions have at most a derivative importance insofar as they figure in the content of this understanding.

But it may be doubted whether intentions have even this derivative role, since the understanding in question can plausibly be interpreted, along the lines sketched in section 3 above, in a way that does not refer directly to intentions. This interpretation would proceed as follows. When we place ourselves in a hospital's care, we do so with the understanding that they will do what they can, according to contemporary medical standards, to treat our illness. We recognize that the competing needs of other patients for scarce medical resources may limit what can be done for us. This is one exception to the duty to provide us with care. But the fact that others would benefit from a worsening of our condition does not constitute the basis for another exception. That is to say, such benefits do not (without our consent) justify hospital staff in knowingly bringing about, or failing to prevent, a worsening in our condition.

If the hospital employees' duties of care have this form, then those who take these duties seriously and comply with them will act with certain intentions. But, as in the analysis of the transplant and terror bombing cases that I offered above, when these duties are violated it is not the fact that the agents have certain intentions that makes these actions wrong. If a doctor were to withhold medicine from me in order to make my organs become available for transplant, what would make this action wrong would not be the fact that the doctor intends that I should die, but rather the fact that what he does violates his duty to treat my illness (a duty to which the need for transplants does not justify an exception.)

I find this a plausible account of the hospital cases. Unlike the explanation of the transplant case offered earlier, however, it does not generalize to cover other cases such as the rescue/transplant example. (There is, of course, no reason why both explanations could not apply in the hospital case: there may be more than one reason why it would be wrong to let one patient die in order to save others by transplant.) An explanation that applied beyond the medical setting would have to be freed from dependence on the conditions of our assent to medical treatment and given a more general moral foundation (such as in the idea of principles that no one could reasonably reject.)

As I suggested above, one such generalization would take the form of a constraint on justifiable exceptions to the principle that prohibits us from acting in ways that it is reasonable to believe will lead to others being killed or injured by the impact on them of our bodies or instruments under our control. The claim would be that it would be reasonable to reject any version of this principle that included exceptions that violate this constraint. A broad constraint of this kind would hold that exceptions to this principle cannot (at least not without the victim's consent) be justified by appeal to the advantages to others of a person's being killed or injured. (As we saw above, self-defense would have to be recognized as an exception to this constraint.) As cases such as the Loop example show, however, there are a number of problems involved in formulating such a constraint and in showing that it could not reasonably be rejected. It may be that any defensible constraint would have to be weaker and more complex than what I have just stated. I cannot explore these problems here.

My aims in this paper have been more limited. I have offered an explanation of some of the cases, such as *Drug/Transplant*, *Rescue/Transplant* and the terror bombing case, in which it has seemed most tempting to explain the wrongness of actions by appealing to the distinction between intended and merely foreseen consequences. My account avoids this distinction, but explains why it might have seemed relevant. This account also suggests a general strategy for understanding such cases: that we should approach them by thinking about the structure of the exceptions that principles restricting harmful conduct can recognize and still be principles that no one could reasonably reject. The general fruitfulness of this strategy remains an open question.