

Draft: Please don't cite without permission

The Road to Hell

Alastair Norcross

1. Introduction: The Doctrine of the Double Effect.

My concern in this paper is a distinction most commonly associated with the Doctrine of the Double Effect (DDE)¹. Elizabeth Anscombe claims that the denial of the DDE "has been the corruption of non-Catholic thought, and its abuse the corruption of Catholic thought."² Many other philosophers, on the other hand, regard the DDE as not just wrong, but also wildly implausible. Some would even say that the most decisive refutation of the DDE is simply a clear statement of that doctrine. I, however, would

¹ There seems to be a schism within the philosophical community concerning the name. For example, Elizabeth Anscombe, Jonathan Bennett and Charlie Curran refer to it as a 'principle' (sometimes adding the definite article before 'double'), while Philippa Foot and Warren Quinn prefer 'doctrine' (with Foot employing the definite article before 'double' and Quinn dropping it). Thomas Nagel, in a truly ecumenical spirit, calls it the 'law of double effect', but also refers to it as both a 'doctrine' and a 'principle'. I choose 'doctrine' (except when I am quoting what others have said), only because I find 'DDE' more catchy than 'PDE' or 'LDE'.

² Elizabeth Anscombe, "War and Murder", from Walter Stein (ed.), *Nuclear Weapons; A Catholic Response*.

The Road to Hell

never go in for such a cheap, though accurate, jibe. Nonetheless, let's begin with some statements of DDE.

Jonathan Bennett:

It permits certain conduct that predictably leads to bad results, if it also leads to good ones and the following are all true:

- (1) The behaviour is not bad in itself.
- (2) The agent's intentions are good.
- (3) The good does not flow from the bad and/or the agent does not intend the bad as a means to the good.
- (4) The good is good enough, compared with the bad, and there is no better route to the former.³

Charlie Curran:

The manuals of theology generally propose the following four conditions under which one can be justified in causing evil in conjunction with good: 1) The action itself is good or indifferent. 2) The good effect and not the evil effect is the one sincerely intended by the agent. 3) The good effect is not produced by means of the evil effect. If the evil effect is not at least equally immediate causally with the good effect, then it becomes a means to the good effect and intended as such. 4) There is a proportionate reason for permitting the foreseen evil effect to occur.⁴

Warren Quinn:

³ Jonathan Bennett, *The Act Itself*, Oxford 1995, pp. 196-7.

⁴ Charles Curran, *Ongoing Revision in Moral Theology*, Notre Dame 1975, pp. 173-4.

The Road to Hell

The doctrine ... is typically put as a set of necessary conditions on morally permissible agency in which a morally questionable bad upshot is foreseen: (a) the intended final end must be good, (b) the intended means to it must be morally acceptable, (c) the foreseen bad upshot must *not* itself be willed (that is, must not be, in some sense, intended), and (d) the good end must be proportionate to the bad upshot (that is, must be important enough to justify the bad upshot).⁵

2. The Means Principle

My concern, and that of most philosophers who discuss DDE, is with the third condition. In fact, some philosophers seem to equate DDE with the third condition. Thus Philippa Foot, in a very influential article: “By ‘the doctrine of double effect’ I mean the thesis that it is sometimes permissible to bring about by oblique intention what one may not directly intend.”⁶ My concern is not with the origins or the interpretation of DDE. Rather, it is simply with the thesis that a particular distinction is morally relevant: the distinction between bringing about an intended bad effect as a means to a good effect and bringing about a foreseen but unintended bad effect in the course of bringing about a good effect. Following Bennett, I will call this *the means principle*. My concern is not

⁵ Warren Quinn, “Actions, intentions, and consequences: The Doctrine of Double Effect”, *Philosophy & Public Affairs* 18:4 (Fall 1989), reprinted in Warren Quinn, *Morality and Action*, Cambridge 1993, p. 175.

⁶ Philippa Foot, “The Problem of Abortion and the Doctrine of the Double Effect”, reprinted in Bonnie Steinbock and Alastair Norcross (eds.) *Killing and Letting Die*, Fordham 1994, p.267.

The Road to Hell

purely theoretical. Some writers appeal to the means principle to argue against certain forms of euthanasia. The United States House of Representatives recently approved legislation that would make it a crime for a doctor to prescribe a drug “for the purpose of causing death”, but would allow the prescription of a controlled substance to alleviate pain, even “if the use of such a substance may increase the risk of death”.

First, a couple of clarifications. Both Bennett’s and Curran’s statements of the third condition refer both to the bad effect being a means to the good effect and to its being intended as such. That is, they refer both to an objective condition and to a subjective condition. Bennett explicitly dismisses the objective condition as “quite implausible: there is no evident reason why a morality should forbid the G[ood]-from-B[ad] causal structure except for what it implies about intentions that may lie behind its production.” (199) I agree with Bennett’s intuitions here, but it is worth considering an example to motivate these intuitions:

Meteor. A meteor is heading towards destroying many people (for those who espouse constraints with thresholds, set the number just below the relevant threshold). We devise a missile that will divert it harmlessly, if launched at the right time. The launch mechanism is risky, so only two qualified scientists, Homer and Barney, are allowed in the launch room. The procedure involves pushing the launch button at noon on the relevant day, and then vacating the room, before the launch takes place. If the button isn’t pressed within ten seconds of noon, the meteor will not be destroyed, and the many will die. Shortly before the scheduled launch, there is an accident in the launch room, which kills Barney and renders Homer unconscious. General Moe, outside the room, sees what has happened. He can’t get in, but he can break a window, which he does.

The Road to Hell

Barney's dead body is slumped in front of the button. Moe has a grenade, which he can toss into the room. The explosion will push Barney's body onto the button, launching the missile, but it will also kill Homer. In this case, Homer's death is not a means to saving the many. Suppose, then, that Moe has acted permissibly, because he hasn't used Homer as a means. But what if he discovers afterwards that the triggering device had been recently fitted with a safety mechanism, which would only allow the launch to proceed if no signs of human life were detected inside the room? Now, it turns out that Moe's action is impermissible, because Homer's death is a means to the saving of the many. But this would make the permissibility of saving many depend on the mechanics of the triggering device. That the mechanics of the triggering device could make a moral difference, without any further difference in terms of who lives or dies, or in terms of Moe's mental states, is too implausible to merit further discussion. Of course, it's not even clear that Curran's statement of the third condition should be read as imputing moral significance to the causal structure itself. Perhaps the claim is that it is wrong for the bad effect to be a means to the good effect, because it will therefore be intended as such. But my example shows that even that cannot be right. The actual causal structure can only determine intentions to the extent that it is known. From now on, I will concentrate on the means principle as a principle about intentions.

A second clarification concerns the imputed force of the means principle. A fairly natural reading of the third condition construes it as a necessary condition on morally permissible action. So, perhaps, we might claim that some things, such as the intentional termination of innocent human life, are absolutely forbidden. Indeed, the infamous

The Road to Hell

statement of the American Medical Association on euthanasia⁷ might be read as appealing to such a prohibition. Even if the good clearly outweighs the bad, and the intended final end is clearly good, we may not intentionally terminate an innocent life. However, the means principle allows us to terminate an innocent life in the pursuit of outweighing good, just so long as the death is not intended as a means to the good end. Elizabeth Anscombe, an ardent supporter of the DDE, sees it as providing exceptions to what would otherwise be absolute prohibitions. "The distinction between the intended, and the merely foreseen, effects of a voluntary action is indeed absolutely essential to Christian ethics. For Christianity forbids a number of things as being bad in themselves. But if I am answerable for the foreseen consequences of an action or refusal, as much as for the action itself, then these prohibitions will break down."⁸ Double effect, says Anscombe, is what rescues the Christian moralist from abandoning her absolutist principles. As I stated the means principle, though, it is quite compatible with the claim that the intentional termination of innocent life is not always wrong. The distinction between bringing about a bad effect that is intended as a means to a good effect and bringing about a bad effect that is foreseen but unintended may make a proportional, rather than lexical, difference, or, even weaker, it may simply make a moral difference when all else is equal. Thus, we could consistently hold that, even if Moe had intended Homer's death as a means to the saving of a billion lives, his action would have been permissible. However, it would have been better, if he had merely foreseen the death. The practical importance of the means principle clearly varies in proportion to its

⁷ For the full statement, see section 4 of this paper.

⁸ Anscombe, *op. cit.*

The Road to Hell

supposed strength. If, as I claim, all versions of the means principle (as applied to actions) are false, it will be nonetheless most important to establish the falsity of the stronger versions, especially because these are the most influential. Therefore, although my concern is with the general thesis, many of my examples will focus on the absolutist versions of prohibitions.

3. Some worries about applying the Means Principle.

So, how are we to apply the means principle? Speaking of a case in which a doctor administers a pain-relieving drug to a mortally ill patient, knowing that the drug may kill the patient if the illness doesn't do so first, Anscombe says: "[E]veryone understands that it is a very different thing so to administer a drug, and to administer it with the intention of killing." But what *is* the difference? It isn't just a matter of desire. Consider a pair of cases often cited in connection with DDE:

Craniotomy: A woman in labor will die unless the head of the fetus she is trying to deliver is crushed. But the fetus may be safely removed if the mother is allowed to die.

Hysterectomy: A pregnant mother's uterus is cancerous and must be removed if she is to be saved. This will kill the fetus. But if no operation is performed the mother will eventually die after giving birth to a healthy infant.⁹

Supporters of DDE often claim that the fetus's death is intended in the craniotomy case, and thus forbidden, but that it is merely foreseen in the hysterectomy case, and thus permissible. In the craniotomy case, though, it is reasonable to assume that the doctor has

⁹ I have taken the descriptions of these cases from Quinn, *op. cit.*

The Road to Hell

no desire that the fetus should die, even though he knows full well that it will. Perhaps it will be objected that the doctor must desire the fetus's death, since it is necessary for something else that the doctor desires, namely the life of the mother. However, the same reasoning could apply to the death of the fetus in the hysterectomy case. If desires were the only criteria of intention, the doctrine of double effect would be in danger of rendering the same judgments as consequentialism about most cases.

Indeed, we could argue that the doctrine can't really distinguish between the craniotomy and the hysterectomy cases. The death of the child in the former case is also an unwanted but foreseen consequence of the surgical procedure and not, strictly speaking, a means to saving the mother's life, given a fairly plausible reading of the doctrine. It is, after all, only an alteration in the shape of the child's head that is required to save the mother's life. While it is true that such an alteration will, with medical certainty, result in the death of the child, it is also true that the hysterectomy will, with equal medical certainty, result in the death of that child. Philippa Foot admits that such an interpretation of the doctrine would "make nonsense of it from the beginning" but insists that there is yet something to it. That something, though, turns out to be a version of the distinction between doing and allowing, and doesn't seem to have anything to do with intentions. This problem threatens to place on the foreseen side of the intended/foreseen distinction all bad effects that intuitively are intended as a means to a good effect.¹⁰ This would clearly deprive the means principle of any force.

¹⁰ For detailed discussions of this issue, see Bennett, *op. Cit.*, pp 201-213, and Quinn, *op. Cit.*, pp. 176-183.

The Road to Hell

A different approach to the means principle involves a different problem. A writer who not only advocates the doctrine of double effect, but also tries to explain the distinction between a foreseen but unintended effect and an effect that is “intended-as-a-means” is John Finnis. In his reply to Judith Thomson's "A Defense of Abortion", Finnis presents four considerations which are supposed to help us decide whether the bad effect of a particular action is to count as intended-as-a-means. Since the context of his discussion is abortion, the bad effect is assumed to be the harming of a victim. "(1) Would the chosen action have been chosen if the victim had not been present?... (2) Is the person making the choice the one whose life is threatened by the presence of the victim?... (3) Does the chosen action involve not merely a denial of aid and succor to someone but an actual intervention that amounts to an assault on the body of that person?... (4) But is the action action against someone who had a duty not to be doing what he is doing, or not to be present where he is present?"¹¹ Finnis himself admits that these four considerations are just factors in the judgment as to whether a particular bad consequence of an act is to count as intended, factors whose relative weights he has not considered. However, they are clearly supposed to be morally relevant factors, for they (or rather (2) (3) and (4)) provide his reason for claiming a moral distinction between Thomson's violinist case and certain abortions. The first thing to note about these considerations is that the last three don't seem to have any necessary connection with the agent's state of mind, which seems a little strange since they are supposed to help us determine whether a particular effect was intended-as-a-means. It seems that Finnis has

¹¹ John Finnis, "The Rights and Wrongs of Abortion", *Philosophy & Public Affairs*, Vol 2 (1973): 117-145.

The Road to Hell

tailored these three considerations to give him a distinction between the violinist case and the case of a life-threatening abortion. I don't think they will be of much use to us in an investigation of intentions, so I shall not discuss them. (1), however, which could be generalized to give a counterfactual test for intention, bears a little close examination.

First, to generalize from the specific context of abortion, we should replace talk of the presence of a victim with that of the occurrence of a harm. Second, since we are concerned with intentions, and not (except tangentially) objective causal structures, we must focus on the agent's *beliefs* about the occurrence of the harm in question. Thus, we get something like the following test: in order to determine whether a particular harm, that resulted from an action, was intended by the agent, we ask whether the agent would have performed the action if s/he had believed the harm wouldn't have occurred. If the answer is yes, the harm is not intended. Lets apply this test to another familiar pair of cases:

Strategic Bomber: A pilot bombs an enemy factory in order to destroy its productive capacity, and thus shorten the war, which results in the saving of over one million innocent lives. However, the pilot foresees that he will certainly kill ten thousand innocent civilians who live near the factory.

Terror Bomber: A pilot deliberately kills ten thousand innocent civilians in order to demoralize the enemy, and thus shorten the war, which results in the saving of over one million innocent lives.

The standard approach to these cases is to claim that the terror bomber intends the deaths that he causes as a means to shortening the war, while the strategic bomber merely foresees the deaths that he causes. Thus, if it is absolutely forbidden intentionally to kill

The Road to Hell

the innocent, the terror bomber acts wrongly. The strategic bomber, on the other hand, may act permissibly, if we judge that saving the lives of over a million outweighs the deaths of ten thousand. Even if there isn't an absolute prohibition against intentionally killing the innocent, the means principle tells us that the behavior of the strategic bomber is morally better than the behavior of the terror bomber. But does the counterfactual test give the result that the terror bomber intends the deaths and the strategic bomber merely foresees them? Here's how it's supposed to go. We ask, of each bomber, would he have dropped the bombs, if he had believed that he wouldn't thereby kill innocent civilians? The terror bomber, we say, clearly wouldn't have dropped the bombs, since he wouldn't then have been able to achieve his objective of demoralizing the enemy and thus shortening the war. The strategic bomber, on the other hand, would have dropped the bombs, and done so more eagerly, since he would have been able to achieve his objective and wouldn't have had to cause the deaths that he, in fact, regrets. But not so fast. It's not clear that this is the right answer. An equally plausible answer is that the strategic bomber wouldn't have dropped his bombs, since, if he had believed that he wouldn't thereby kill innocent civilians, given what he knew about the proximity of the civilians to the factory, he would have believed either that he was off target and that he wouldn't have hit the factory, or that his bombs would not have been powerful enough to destroy the factory. Likewise, in the hysterectomy case, if we asked whether the doctor would have performed the operation, if she had believed that she wouldn't thereby have killed the fetus, we could just as easily answer no as yes. Given what the doctor knew about the patient (for example, that she was pregnant), if she believed that a hysterectomy wouldn't

The Road to Hell

have resulted in the death of a fetus, she would have believed that she was operating on a different patient.

We could avoid this result by specifying that the only change in the agent's beliefs be that the harm doesn't occur. The problem then, of course, is that we get the wrong result in the case of the terror bomber. If the *only* change in the terror bomber's beliefs were the elimination of his belief that civilians would be killed by his bombs, he would still have bombed (and done so gladly), because his belief that the bombing raid would demoralize the enemy would have remained. It seems as though we will either count too much or too little. Jonathan Bennett suggests, in response to this dilemma, that we specify that the change in the agent's beliefs be restricted to the belief that the harm won't occur and to any changes in her belief set that follow from that change by a "causally downstream inference".¹² So the strategic bomber would still have carried out the raid, because the changes in his beliefs that I suggested—that he was either off course or had inadequate bombs—were "causally upstream", that is, they involved thinking about what would have had to have been different prior to dropping the bombs in order for it to be true that the bombs wouldn't have killed civilians. The terror bomber, on the other hand, would not have carried out the raid, because his changed belief about civilian deaths would also require changed beliefs about what would or wouldn't have followed. In particular, he would have believed that the enemy wouldn't have been demoralized, and thus wouldn't have surrendered. Although Bennett claims that his suggestion correctly classifies the terror and strategic bombers as being on opposite sides of the intending/foreseeing line, he doesn't claim that this distinction has any moral

¹² Bennett, Op. Cit. P. 217

The Road to Hell

significance. Of the two bombers, he says, “To get them apart we had to specify what causally follows downstream as distinct from what causally follows upstream. Why should that difference matter morally?”¹³

I agree that Bennett’s suggested distinction has no moral significance. It is not so clear, though, that Bennett’s account, in terms of the distinction between causally upstream and causally downstream inferences, correctly captures the intending/foreseeing distinction. Recall that the strategic bomber is said to foresee the civilian deaths, because he would still have carried out the raid if he had believed that the deaths wouldn’t have occurred, and had believed whatever followed by a “causally downstream inference”. He would still have believed that the factory would have been destroyed, and thus that the war would be shortened. But does this necessarily follow? Given his beliefs about the destructive capacity of the bombs and the proximity of the civilians—beliefs which are causally upstream from his changed belief about the survival of the civilians—couldn’t he have reasoned that the unexpected survival of the civilians would have had detrimental effects on the course of the war? For example, his side might have become demoralized at the apparent indestructibility of the enemy civilians, believing that the other side had developed advanced defensive capabilities. Or perhaps the enemy’s resolve would have been strengthened, in the belief that god had miraculously saved them from certain destruction. After all, how else could they explain the fact that bombs powerful enough to destroy factories had no effect on civilians who were right next to them? Such alterations in the bomber’s belief set, clearly causally downstream from the belief that no civilians would be harmed by the bombs, could well result in the bomber canceling the

¹³ Ibid.

The Road to Hell

raid. Do such counterfactual speculations incline us, at all, to the view that the strategic bomber intends the deaths of the civilians in the actual world? Of course, in each case, we could specify exactly how the counterfactual should be read to achieve the intuitively acceptable results, but this sort of ad hoc gerrymandering should make us suspicious. Perhaps the whole notion of intention, at least as applied to means, is too fragile to do any useful moral work.

4. Against the Means Principle.

The worries I have sketched about even applying intentions in morally useful principles deserve more extensive treatment than I can give them here. They may yet prove fatal. However, I am going to put them aside for the remainder of this paper, and proceed on the assumption that there is a philosophically satisfactory distinction between bringing about an intended bad effect as a means to a good effect and bringing about a foreseen but unintended bad effect in the course of bringing about a good effect. Furthermore, I will assume that the standard cases correctly illustrate the distinction. After all, part of our dissatisfaction with the different variants of the counterfactual account stems from their inability to classify the standard cases in accordance with our intuitive judgments of them. Consider again the terror bomber's raid and the strategic bomber's raid. It is often claimed that the DDE renders the latter permissible but the former impermissible. Now consider a variation on the terror bomber. The bomber can demoralize the enemy by killing one particular man, say the famous sporting hero Schultz. If the intentional killing of the innocent is absolutely forbidden, it is impermissible to save a million lives by killing Schultz. But now, let's say, the very same bomber can perform the strategic raid

The Road to Hell

in which ten thousand innocent civilians are killed. Furthermore, the deaths of the ten thousand would be much slower and more painful than Shultz's death in the terror raid, though not so bad as to outweigh the saving of a million lives. And, to crown it all, Shultz is one of the ten thousand (imagine that the terror raid can be carried out on Shultz with a very accurate missile during the daytime, when he is away from the town, hiking, but the strategic raid can only be carried out with inaccurate missiles at night, when Shultz is back in his house by the factory). The absolutist version of the DDE tells us that Shultz's instant painless death in the terror raid is the intended means to the saving of the million lives, and thus impermissible, but his much slower and more painful death, and those of 9,999 innocent others, in the strategic raid is the unintended, but foreseen, side-effect of the means to the saving of the million lives, and thus permissible. That is, if our bomber is faced with this choice between killing Schultz quickly and painlessly and killing him and 9,999 others slowly and painfully, he must choose the latter. The moral bankruptcy of this result needs no further comment. If we apply the means principle to a non-absolutist morality, we might not be able to get such an extreme contrast, but we will be able to construct clearly implausible results along similar lines (I leave that as an exercise to the reader).

As I said earlier, the means principle has been invoked in the context of euthanasia. The AMA policy statement on euthanasia, adopted by the House of Delegates of the American Medical Association, December 4, 1973, forbids the 'intentional termination of life'. Here is the statement:

The intentional termination of the life of one human being by another—mercy killing—is contrary to that for which the medical profession stands and is contrary to

The Road to Hell

the policy of the American Medical Association.

The cessation of the employment of extraordinary means to prolong the life of the body when there is irrefutable evidence that biological death is imminent is the decision of the patient and/or his immediate family. The advice and judgment of the physician should be freely available to the patient and/or his immediate family.

James Rachels, in an influential and widely reprinted article¹⁴, criticizes the AMA policy for drawing an illegitimate distinction between killing and letting die. Bonnie Steinbock and Thomas Sullivan respond that the relevant distinction is between intending and foreseeing¹⁵. Against Rachels' claim that the cessation of treatment in cases where death is known to ensue simply is the intentional termination of life, both Steinbock and Sullivan argue that life-support may be terminated without the intention to terminate life, although such a termination will be foreseen:

But we all know that it is entirely possible that the unwillingness of a physician to use extraordinary means for preserving life may be prompted not by a determination to bring about death, but by other motives. For example, he may realize that further treatment may offer little hope of reversing the dying process and/or be excruciating, as in the case when a massively necrotic bowel condition in a neonate is out of control. The doctor who does what he can to comfort the infant but does not submit it

¹⁴ James Rachels, "Active and Passive Euthanasia", in Steinbock and Norcross, op. cit.

¹⁵ Bonnie Steinbock, "The Intentional Termination of Life", Thomas Sullivan, "Active and Passive Euthanasia: An Impertinent Distinction?", both in Steinbock and Norcross, op. cit.

The Road to Hell

to further treatment or surgery may foresee that the decision will hasten death, but it certainly doesn't follow from that fact that he intends to bring about its death.¹⁶

In response, Rachels considers two doctors, Doctor White and Doctor Black, each faced with the necrotic neonate of Sullivan's example. Doctor White ceases treatment because he regards it as pointless, and thinks that it will only increase the suffering. He foresees that death will occur, but he doesn't "seek, choose, or plan that death, so it is not part of his intention that the baby dies."¹⁷ Doctor Black also realizes that further treatment is pointless and will only increase the infant's suffering. But his intentions are different: "He decides that it is better for the baby to die a bit sooner than to go on suffering pointlessly; so, with the intention of letting the baby die, he ceases treatment."¹⁸

According to Sullivan's interpretation of the AMA doctrine, Doctor Black's action is impermissible, because it is the intentional termination of the infant's life. Doctor White's action, on the other hand, appears to be permitted by the doctrine. But their actions cannot differ in permissibility, says Rachels, since they did "*the very same thing*". More generally, Rachels claims that intentions never affect the moral status of actions, but only of the character of agents.

Rachels' argument that the actions of Black and White could not differ morally seems to depend on a substantive and controversial account of action-type individuation. It is because they do the "very same thing" (i.e., the same type of thing) that the morality

¹⁶ Sullivan, op. cit., 135-6.

¹⁷ James Rachels, "More Impertinent Distinctions and a Defense of Active Euthanasia", in Steinbock and Norcross, op. cit., 141.

¹⁸ Rachels, op. cit., 142.

The Road to Hell

of their actions cannot differ morally. Even if we grant, as the argument implies we should, that two tokens of the same action type cannot differ morally, must we also grant that the actions of Black and White are of the same type? Could there not be a plausible theory of action-type individuation according to which the intention with which an action is performed is relevant to its type?¹⁹ I won't explore this issue further here, however, since it seems to me that Rachels' example itself provides a powerful reason for denying that intentions are relevant to the moral status of actions. Even if we grant that Doctor Black and Doctor White do different types of things, how could we possibly claim that one is permissible but the other isn't? Look at it from the point of view of the infants themselves. Are we going to refuse Doctor Black's patient relief from suffering, because it had the bad luck to get stuck with a doctor who aimed at its death rather than merely foresaw its death? When she complains (metaphorically) that the patient in the next bed, who seems to be in the same condition, is being treated differently, and more humanely, and asks why Doctor Black isn't allowed to do the same for her, we'll explain that her neighbor has Doctor White for a neonatologist, and that what would appear to be the same treatment from Doctor Black would actually be different, because of certain of Doctor Black's mental states. This difference in the mental states of the two doctors, we continue, explains why she must suffer far more excruciating pain than her neighbor.

The example of Doctors Black and White is bad enough, but at least it has few serious practical implications. Even if the difference in the doctors' mental states does render one act permissible and the other impermissible, there will be problems enforcing

¹⁹ For an argument along these lines, see Thomas Sullivan, "Coming to Terms: A Response to Rachels", in Steinbock and Norcross.

The Road to Hell

a policy based on this distinction. Doctor Black will be aware of the policy. If he really wants his patient to die, he will simulate the intentions of Doctor White. How are we to know that his decision to stop treatment is really aimed at the death of his patient, if he doesn't tell us? A suitably designed process of moral indoctrination may even succeed in persuading Doctor Black that it is wrong to aim at Patients' deaths, and thus perhaps change his intentions. But there is more trouble in the neighborhood. Earlier I quoted Anscombe on the distinction between administering an overdose of painkillers with the intent to kill and with the intent to relieve pain but foreseeing certain death. Here is Arthur Dyck on a similar topic:

From the point of view of the dying person, when could his or her decisions be called a deliberate act to end life, the act we usually describe as suicide? Only, it seems to me, when the dying person commits an act that has the immediate intent of ending life and has no other purpose. That act may be to use, or ask the physician to use, a chemical or an instrument that has no other immediate effect than to end the dying person's life. If, for the sake of relieving pain, a dying person chooses drugs administered in potent doses, the intent of this act is not to shorten life, even though it has that effect.²⁰

(Strictly speaking, Dyck's claims entail that very few, if any, perform 'a deliberate act to end life'. Surely, the act of terminating a suffering life has the purpose of ending suffering, as well as ending life. A charitable reading of Dyck will omit the careless phrase 'and has no other purpose'.) Both Anscombe and Dyck claim that it is possible to

²⁰ Arthur J. Dyck, "An Alternative to the Ethic of Euthanasia", reprinted in James E.

White (ed.) *Contemporary Moral Problems*, 5th ed. West 1997, pp. 212-216

The Road to Hell

administer overdoses of painkillers without the intent to kill, even though the death is foreseen with certainty. Likewise it is possible to administer such an overdose with the intent to kill. We could consider a version of the Doctors Black and White example based on this difference, but I'll leave that to you. My concern is elsewhere.

Even though overdoses of painkillers can be knowingly administered without the intent to kill, presumably there are some forms of causing death that cannot be so employed. Anscombe says "It is nonsense to pretend that you do not intend to do what is the means you take to your chosen end."²¹ Dyck, in the passage quoted above, talks of drugs that have "no other immediate effect than to end the dying person's life". An overdose of painkillers can be administered without lethal intention, because painkillers have the effect of relieving pain as well as (in large enough doses) ending life. But what about a dose of some toxin that causes instant painless death, but doesn't directly work on the pain centers? Such a substance could not, according to Anscombe and Dyck, be knowingly administered without deadly intent. (If it could, that would make a nonsense of their position.) Imagine now two patients, Lucky and Loser, both terminally ill and in unrelievable (except by death) excruciating agony. Both beg their doctor to administer a lethal dose of something, anything, to end their suffering. Lucky's doctor has a suitable painkiller that will, in a large enough dose, produce instant painless death. Loser's doctor has neglected to replenish his supply of painkillers, but he does have a poison that will also produce instant painless death. Phenomenologically, the experiences of Lucky and Loser would be identical. However, the lethal substances produce death by slightly different mechanisms, and so, knowing this, Loser's doctor cannot administer the poison

²¹ Anscombe, *op. cit.*

The Road to Hell

without lethal intent. Loser is thus condemned to an agonized death, while Lucky gets his wish of a painless death. I won't dwell on the moral horror of this conclusion, because there is worse to come.

Imagine now that a new painkiller is discovered that is far superior to all existing painkillers. Call it *SuperMorph*. SuperMorph works superbly on all patients, with no side effects and no problems of addiction. Even better, SuperMorph is cheap and easy to produce. The medical establishment uses it exclusively and discontinues production of all other painkillers. There are still, however, a few cases of terminal pain that can only be relieved by a lethal dose of SuperMorph. In which case, death is instant and painless. When a patient reaches this stage of a terminal disease and requests an overdose of SuperMorph, her doctor administers it with the purest intention of relieving suffering, though she also foresees death with certainty. Consider now the cases of Lucky and Loser again. This time both their doctors have plenty of SuperMorph. However, Loser's doctor, Doctor Smart, has made a discovery about SuperMorph. Even though SuperMorph works just like a regular painkiller (but better) in nonlethal doses, it works completely differently in lethal doses. In fact, a lethal dose of SuperMorph works exactly like the poison that Loser's doctor had in the previous example. It doesn't work directly on the paincenters at all. Thus, a lethal dose of SuperMorph involves causing the death of the patient as a means to relieving suffering. The research that led to this discovery is extremely complicated and hard to follow, though. In fact, it is too hard to follow for anybody except Doctor Smart, who is, by far, the smartest man who ever lived. Unfortunately (though not for most terminal patients), Smart is also extremely erratic and eccentric, with a track record of wildly implausible claims. Nobody, therefore, believes

The Road to Hell

his claims about the workings of lethal doses of SuperMorph, and they are perfectly justified in their disbelief. Lucky's doctor, Doctor Average, is pretty smart himself, but he believes, along with the rest of the world except Smart, that overdoses of SuperMorph produce death as a by-product of relieving pain. He thus administers the requested dose to Lucky with perfectly permissible intent. Doctor Smart, however, may not administer the requested dose to Loser, since he would have deadly intent if he did so. What is more, Loser doesn't have access to any other physicians. Thus, Loser is again condemned to an agonized death. This time, though, the only difference between Lucky's and Loser's positions is that Loser has an exceptionally smart doctor and Lucky has a merely normally smart doctor. Clearly, this cannot be the right result. The means principle has led us horribly astray.

5. The Means Principle applied to character.

What, then, is the moral significance of intentions, if any? I suggest, in common with other consequentialist moralists, that the moral significance of intentions concerns our appraisals of agents as opposed to actions. The primary source for our judgments of agents is, of course, their actions, but intentions are significant too. Consider the following case: the President of the US is at a public rally surrounded by crowds of people. A security man suddenly raises his gun and shoots a would-be assassin who was within a few feet of the President with his own gun raised. The security man is praised for saving the President's life. If we were to discover that the security man was totally unaware of the presence of the assassin and had, in fact, tried to shoot the President, but by a lucky chance had not only missed but actually hit the other killer, we would, I think,

The Road to Hell

judge him somewhat differently. We would, of course, still be thankful that he had shot, for if he hadn't, the President would have been killed by the other man²². We would not have altered our judgment of the action, but only of the character of the security man. Consider again the case of the fetus-destroying hysterectomy. Let us say we are quite satisfied that it is the right thing to do in the circumstances. We regret the death of the fetus, but we value the life of the mother more highly. Now suppose we discover that the doctor who performs the operation had a burning desire to kill the fetus, and that was the main reason he agreed to perform the operation. This doctor gets a lot of personal satisfaction out of killing fetuses. We would probably be horrified to learn this and would think the doctor utterly reprehensible, but would we think that the operation should not have been performed (suppose there were no other doctors available who could have done it in time)?

Notice, though, that my examples concern what is intended as an end. If we add to the presidential assassination case the detail that the security man has recently discovered that the president plans to abolish the minimum wage and the Department of Education, withdraw the United States from the United Nations, return the dollar to the gold standard, and declare homosexuality to be contrary to god's law²³, our intuitive judgements of his character may change. In this case, the security man would not intend the death of the president as an end in itself, but rather as a means to preventing further

²² Lest I seem to be exaggerating our thankfulness at the president's survival, imagine that the president is Josiah Bartlett, from the TV drama "The West Wing".

²³ All items on the Texas Republican Party's official platform, adopted at its 2000 convention.

The Road to Hell

moral horror. If the president's plans don't seem sufficiently horrific, add his plan to exterminate all US residents who are not fundamentalist protestant Christians²⁴. Is there, then, anything to be said for a version of the means principle that applies to character?

Consider again the strategic and terror bombers of the original example. Both bombers knowingly cause the deaths of ten thousand innocent civilians in the course of saving a million lives. The strategic bomber doesn't intend these deaths as a means to his goal, but he does foresee them with certainty. The terror bomber intends the deaths as a means to his goal. Assume further that the strategic bomber would refuse to perform the terror raid, but that the terror bomber would be equally willing to perform the strategic raid, were it an option. Is there any reason to think, given these details, that the terror bomber has a worse character than the strategic bomber?

Assume, for the moment, that character is, in some sense, a measure of one's propensity to perform better or worse actions. This approach can be roughly characterized as follows: One character trait, C1 is better than another, C2, just in case the possession of C1 makes one likely, *ceteris paribus*, to perform better actions than does the possession of C2. This approach can be subject to many variations. For example, do we compare C1 and C2 with respect to a particular person, a particular type of person, the "average" person, etc.? Do we compare propensities with respect to the circumstances a particular individual is likely to encounter, given what we know about her, given her social position, given "normal" circumstances, etc.? It may be, for example, that certain combinations of character traits that would be bad in most people

²⁴ Clearly we are not talking about Josiah Bartlett any more. This plan, surprisingly, didn't even make it to the official platform of the Texas Republican party.

The Road to Hell

would be good in politicians, or soldiers, or hockey players. This may make it appropriate to describe someone as a good politician but a bad person. What we mean by that (or rather what we ought to mean by that) is that, given that she is a politician, she is a good person, but had she had the same character but a different calling, she would have been a bad person.

Back to the strategic and terror bombers. If my earlier argument against the means principle as applied to actions is correct, there is an obvious respect in which the terror bomber has a *better* character than the strategic bomber. Given certain choices between two different war-ending raids, the strategic bomber would perform the worse action, but the terror bomber would perform the better one. Furthermore, if the only way to end the war were to carry out the terror raid, the strategic bomber would refuse, thus condemning a million people to death (if there were no other suitable pilot available). Are there any situations in which the strategic bomber would make a better choice than the terror bomber? Clearly there are not, if they are sincere, accurate, and not self-deceived in their assessments of the harms and benefits that result from their actions. However, it might be argued, insincerity, inaccuracy, and self-deception are common human traits. Given this, the willingness to bring about harm as a means to a greater benefit is likely to result in actions that bring about harm as a means to a lesser benefit. Maybe so. But there are two lines of response to this point that I will briefly pursue.

First, the same point can be made with respect to the willingness to cause harm as an unintended but foreseen effect of an action that also causes a greater benefit. In fact, there is at least intuitive reason to think that this willingness is more liable to abuse than the willingness to cause harm as a means to greater benefit. To the extent that harms are

The Road to Hell

intended as means to one's ultimate goal of greater benefit, they are more likely to be vividly represented in the imagination than are harms that are unintended but foreseen side-effects. To the extent that they are more vividly imagined, they are less likely to be underestimated in one's eagerness to bring about the perceived greater benefit. To the extent that either willingness is a bad character trait, then, it would appear that the willingness to bring about harm as a foreseen but unintended effect of an action that also causes a greater benefit is actually *worse* than the willingness to bring about harm as a means to a greater benefit. It is entirely possible that we should teach people to accept harm to others *only* when it is intended as a means to a greater good, and not when it is a "merely" foreseen side-effect. This suggestion would be almost a complete reversal of conventional thinking, especially as applied to military matters. That is hardly a reason to reject it, though.

Second, if we are to count the possible abuses of certain character traits as evidence of their undesirability, we must weigh these against the possible and actual bad effects of possessing various proposed alternative character traits. Just as the willingness to cause harm as a means to a greater benefit may lead to bad choices, so may the unwillingness, either absolute or conditional, to do so. In fact, there are many actual cases in which such an unwillingness does lead to very bad choices, most notably those involving a refusal to comply with dying patients' rational requests to end their lives on their own terms. It is highly unlikely that any attempted calculation, taking reasonably into account what we know of human psychology and the kinds of significant choices facing us, would come down firmly on the side of an unwillingness to cause harm as a means to a greater benefit. Questions of which character traits would lead to which

The Road to Hell

results are, of course, empirical ones. I am not claiming that a character analog of the means principle *could* not be justified. I am claiming that, in advance of extensive psychological investigation, there is no more intuitive reason to think that such a character trait will be justified than that a more utilitarian character trait will be justified.

Of course, my whole approach to assessments of character traits might be challenged, on the grounds that it illegitimately ties assessment of character to assessment of action in such a way that the latter has theoretical primacy. Certain versions of virtue ethics either divorce assessments of character from assessments of actions altogether, or give theoretical primacy to the former. According to this criticism, the question of whether a character analog of the means principle is true has nothing to do with what actions are likely to issue from different kinds of willingness to bring about harm. I find such approaches to character deeply mystifying, but, to the extent that I understand them, I can see no reason to prefer the willingness to cause harm as a foreseen but unintended effect of an action that also causes a greater good over the willingness to cause harm as a means to a greater good. I will, however, briefly consider one line of argument to the contrary.

Consider the kantian prohibition on using a rational agent as a mere means. Such a prohibition is, in kantian thought, closely tied to the notion of respect for persons. It might then be argued, on roughly kantian lines, that a willingness to cause harm to a person as a means to a greater good involves a lack of morally appropriate respect. In order to provide an argument for a character analog of the means principle, it would also have to be argued that a willingness to bring about harm to a person as a foreseen but unintended effect of an action that also causes a greater good does not involve such a lack

The Road to Hell

of respect. Perhaps it involves no lack of respect, or perhaps it involves a less serious lack of respect. But how would such an argument go? It is not in the least clear to me that the willingness to cause me harm as a side-effect of an intended means to a good end involves any more respect for me than does the willingness to cause me harm as an intended means to the same good end. Either willingness is compatible with a recognition of my moral worth. Neither willingness necessarily involves a failure to take my interests into account. It is precisely because the utilitarian recognizes the equal moral worth of everyone that she is prepared to harm me as a means to a greater good. A refusal to do so seems to involve a failure to respect the others who could be helped by harming me. It is true that many people profess a dislike of being used. The usual contexts in which people consider the matter, though, involve being used to further someone else's selfish goals. Furthermore, such contexts typically involve the user making little or no attempt to take the victim's interests into account. If I were told that I could either be harmed as a means to bringing about a greater good, or harmed to the same degree as a foreseen but unintended side-effect of bringing about the same greater good, it's not clear that it would be rational to prefer one over the other. Whether or not it would be rational, it would be psychologically understandable to prefer the *former* over the latter. After all, if the harm to me were actually a means to producing the good effect, there's a sense in which my suffering (or even death) would have meaning. I could say to myself, "at least my death achieves something good". If my death were simply a side-effect of the means to the good effect, it could seem as if my death were insignificant, and perhaps even pointless. I can see no reason, then, to accept this kantian²⁵ argument

²⁵ My use of the term 'kantian' here and elsewhere should not be taken to imply anything

The Road to Hell

for a character analog of the means principle, save for a dogmatic insistence that the willingness to harm someone as a means to a greater good involves a lack of morally appropriate respect.

If intentions are, as I suggest, only relevant to our judgments concerning moral character²⁶, and even then only final intentions, they cannot enter into a principle which distinguishes certain acts as morally permissible and others as impermissible.

Furthermore, it is far from clear that they can enter into a character analog of the means principle. Truly, then, the road to Hell is paved with bad distinctions.

University of Colorado Boulder

about whether Kant himself would have agreed with any aspect of the argument.

²⁶ For an interesting argument that suggests the contrary, see Steven Sverdlik, "Motive and Rightness", *Ethics*, Vol. 106 No. 2 (1996): 327-349.