

Intending, Foreseeing and the Doctrine of Double Effect

by

Ann Bumpus

Submitted to the Department of Linguistics and
Philosophy in Partial Fulfillment of the
Requirements for the Degree of

Doctor of Philosophy

in

Philosophy

at the

Massachusetts Institute of Technology

February 1995

© 1995 Ann Bumpus

The author hereby grants to M.I.T. permission to reproduce and to
distribute publicly paper and electronic copies of this thesis
document in whole or in part.

Signature of Author _____

Department of Linguistics and Philosophy
December 7, 1994

Certified by _____

Judith Jarvis Thomson
Thesis Supervisor

Accepted by _____

Ned Block
Chair, Departmental Graduate Committee

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

INTENDING, FORESEEING AND THE DOCTRINE OF DOUBLE EFFECT

by

ANN BUMPUS

Submitted to the Department of Linguistics and Philosophy on December 9, 1994 in partial fulfillment of the requirements for the Degree of Doctor of Philosophy in Philosophy

ABSTRACT

We typically assume that there is a difference between foreseeing an effect of one's voluntary action and intending the effect. Call the view that there is such a difference 'the Ordinary View'. My dissertation is a defense of the Ordinary View against two recent challenges.

The first challenge to the Ordinary View I call "Holism". The upshot of the holist's position is that we intend all the foreseen effects of our voluntary actions. I begin by considering and arguing against a couple of different versions of holism.

The second challenge to the Ordinary View is made by Jonathan Bennett. Bennett argues that there is no unproblematic way to mark off effects of action which are intended as a means from effects which are merely foreseen. His conclusion is that the notion 'intended as a means' is incoherent.

These challenges to the Ordinary View are of interest because they threaten to undermine the Doctrine of Double Effect which claims that whether an agent intends or merely foresees a harmful effect of action sometimes matters to the moral permissibility of performing the action. In response to Bennett, I first examine a view put forth by Warren Quinn. Quinn concedes to Bennett that the distinction between intended versus merely foreseen effects faces serious problems, but attempts to limit the ethical implications of Bennett's challenge by offering a replacement principle for DDE which he believes is not subject to Bennett's objection. I defend Quinn's replacement principle against certain recent attacks in the literature, but point out that Quinn's replacement principle really doesn't escape Bennett's objection. I then return to Bennett's objection and argue that it rests upon two false assumptions. I conclude that the Ordinary View withstands the two challenges.

Thesis Supervisor: Judith Jarvis Thomson

Title: Laurance S. Rockefeller Professor of Philosophy

Acknowledgements

I would like to thank:

Judy Thomson for her insightful comments all along the way, for providing a model of philosophical clarity and rigor, for her remarkable ability as an advisor to ask just the right question, and not least, for her endurance.

David Brink for providing lengthy and detailed helpful comments on numerous drafts as well as much encouragement throughout. I am especially grateful for his help at the early stages of defining the project.

Michael Bratman for his comments on the material in chapters 3 and 4 and for running a terrific NEH Seminar on intention.

Joshua Cohen and Dick Cartwright for joining my thesis committee at the last minute and providing helpful comments.

Martin Davies with whom I have discussed much of the dissertation.

All the members of the philosophy department at Dartmouth College where I have presented much of this thesis. They have provided numerous helpful suggestions on much of this material as well as general encouragement and moral support.

The members of the 1993 NEH Seminar on Intention. I am grateful to all participants of the seminar for their comments on my work and their thought provoking discussions on the topic of intention. The members of the seminar were: Gary Fuller, Eric Gampel, Phil Gasper, Steward Goetz, Thomas Kapitan, Tommy Lott, Al Mele, Tika Newton, Piers Rawling, Harry Silverstein, Steve Sverdlik. Special thanks to Michael Bratman and Eric Gampel for many hours of discussion about the latter chapters.

George Boolos for his ongoing friendship, support and encouragement.

And finally, Michael Anthony, David Galloway, Eric Lormand, Paul Pietroski, Gabriel Segal and especially to my good friends and office mates Susan Dwyer, Jim Page, Sarah Patterson, and Trafford Welden, who helped in the indirect but important way of making graduate school much more fun than it would have been without them.

CONTENTS

I. The Holistic Challenge	9
II. Bennett's Challenge	34
III. Direct Agency: Quinn's Response to Bennett	58
IV. Refutation of Bennett's Challenge	97
V. Intending as a Means	113
VI. Conclusion	120

Chapter One: The Holistic Challenge

1. Introduction

Consider the cases of two bombers fighting for their country:

STRATEGIC BOMBER: A strategic bomber aims to weaken the enemy militarily by bombing a key weapon factory. He sees that the factory cannot be destroyed without causing the deaths of some nearby civilians. Although he would prefer not to cause civilians deaths, he decides to go ahead and bomb.

TERROR BOMBER: A terror bomber aims to weaken the enemy's morale by killing civilians. He targets a civilian shelter which happens to be located near a weapon factory. He foresees that by bombing the shelter, he will bring about the factory's destruction as well, but he's indifferent to this latter effect. He decides to bomb.

Assume each bomber acts on his decision to bomb. Then, under an assumption about the individuation of actions, each of the bombers does the same thing--each one drops a bomb and brings about the destruction of a weapons factory and the deaths of some children. Despite this, some people

think that what the terror bomber does is morally worse than what the strategic bomber does. Why? People who believe the cases are morally different may have the following in mind: the strategic bomber foresees that by bombing he will bring about the deaths of the civilians, but he does not intend to bring about the deaths. The terror bomber, on the other hand, does intend to bring about the civilian deaths; this is his aim. So these people may think it matters morally whether a person intends to bring about some harm or merely foresees that he will do so. This view stands in opposition to consequentialist views since it draws a moral distinction between two actions despite the fact that the actions have exactly the same consequences.

The bomber cases are meant to be illustrative of a moral principle according to which the strategic bomber's action is permissible and the terror bomber's is not. The principle, called the 'Doctrine of Double Effect' is found in Catholic theology and is usually traced back to Aquinas. Elizabeth Anscombe brought the principle into secular philosophy in her essay *War and Murder*.¹ The doctrine is this:

The Doctrine of Double Effect (DDE): Where an agent, X, expects to produce two effects by A-ing, one of them bad and the other good, it is impermissible for X to A if she intends to bring about the bad effect and permissible for her to A if she does not intend to bring about the bad effect.²

¹See St. Thomas Aquinas, *Summa Theologica*, II-II, p.64,a.7.c. Also see G.E.M. Anscombe "War and Murder" (1981).

²On most views, what I call 'DDE' requires qualification: if the harm the agent would bring about as a foreseen but unintended effect of A-ing is a grave enough harm, then the agent may not proceed. The following may be such a case: the strategic bomber has a very powerful bomb and foresees that by bombing the armory he will destroy an entire city. Although he does not intend this result, the harm is too serious for the action to be permitted by the DDE. I ignore the need for qualification here; it is not relevant for our purposes.

DDE assumes (1) that there is a difference between intending to bring about an effect and foreseeing you will do so and (2) that it is morally worse to cause harm which you intend than it is to cause harm which you foresee. In this chapter and the next, I will not be concerned with the second claim, or with any other moral claim. Rather I will be interested in the first claim, which I will call 'the Ordinary View'.

The Ordinary View: There is a difference between intending to bring about some effect by A-ing and merely foreseeing that by A-ing you will bring about the effect.

Take a very common sort of occurrence. You are tired but you want to get some work done. You decide to drink some coffee, knowing that coffee makes you jumpy but also makes you more alert. Ordinarily we would say that you intend to drink the coffee and that you intend to bring about your increased alertness but you do not intend to bring about the jumpiness. We sometimes make this distinction by calling what is unintended but foreseen a "side effect", so in this case, we would say you bring about jumpiness as a side effect of drinking the coffee; it is not an effect you intend to bring about. If some foreseen effects of action are side-effects, then the Ordinary View is true. On the other hand, if an agent intends to bring about all the effects she foresees, then DDE is empty. As plausible as the Ordinary View seems, it has come under attack and I am interested in whether or not it can be defended. In this chapter, I will present and argue against two positions which challenge the Ordinary View: Broad Holism and Deliberative Holism.

2. Broad Holism

In describing the bomber cases, I said that each bomber decides to bomb. Someone who holds a holistic view of decision-making would argue that the bomber does not decide simply to bomb; rather the bomber decides on something broader, he decides on a scenario. A holist about decisions would point out that each bomber deliberates over whether to destroy the munitions factory and kill the civilians, and so on. The holist would then argue that when the bomber reaches a decision, he decides on one of these broader packages.

A simple analogy may help to make this holist's point more plausible. Imagine yourself at the market trying to decide between the Lean Cuisine Chicken Burrito Dinner and the shrink-wrapped New England Boiled Dinner. Both are, of course, extremely tempting. In the end your love of cabbage prevails and so you choose... well, what do you choose? It seems most accurate to say you choose the New England Boiled Dinner, not simply the cabbage. You choose this dinner because it contains the cabbage, but still you choose the whole package over the Burrito package.

The holist is making a similar point about decisions to act. The strategic bomber does not simply decide to destroy the armory. Rather he decides on a whole package; the package includes destroying the armory but it also includes bringing about civilian deaths. I will call these packages which agents allegedly decide upon 'scenarios'. I will take scenarios to consist only of actions. The strategic bomber might be said to compare a scenario which includes destroying the armory and killing the civilians and weakening the enemy to a second scenario which includes his returning to the base (without having bombed) and explaining to a commanding officer why he did not go through with the raid. In discussing scenarios, I will often simplify in the

following two ways. First, though I do not always make this explicit, I assume that the bomber is considering *his* killing the civilians, etc. So if each of two agents, X and Y, are considering whether to A and to B, they are not thereby considering the same scenario. X is considering a scenario which includes X's A-ing and X's B-ing while Y is considering a scenario which includes Y's A-ing and Y's B-ing.

Second, I will often refer to the scenario in which an agent A's (the scenario which includes the strategic bomber's bombing for example). More realistically, the strategic bomber might be said to be considering many scenarios which include his bombing. He might, for example, be uncertain about how the world will be if he bombs. Perhaps he is uncertain about whether a nearby church will be destroyed as well. In this case, he may consider two scenarios, one in which he bombs but does not destroy the church and another in which he bombs and does destroy the church. (Then presumably he would make some assignment of probabilities to these different scenarios.) I am going to ignore such complications.

A further question arises about the breadth of scenarios. For example, does a scenario in which an agent A's include her bringing about all foreseen effects of A-ing? Here there will be a certain looseness in how I use the term 'scenario'. I want it to stand in for whatever these packages are which holists claim people deliberate over and decide upon. But I will be presenting two different views of holism, and these views will differ on how broad these packages are, thus on one view (Broad Holism), scenarios will be broader than on the other view (Deliberative Holism). The Broad Holist's point about decisions may be summarized as follows:

Broad Holistic Decision Principle: If X foresees that A-ing will result in e and decides to A, then she decides on a scenario which includes A-ing *and* bringing about e.

Notice that on this view, an agent who decides to A, decides on a scenario which includes not only her A-ing but her bringing about all foreseen consequences of A-ing. This is the broad holistic view of decisions; we don't simply decide to A, but decide on these broader packages (just as we decide not only on the cabbage, but on the whole dinner).

So far we have been looking at a claim about decisions. How does all this bear on intentions? The following principle expresses a very plausible view about the relation between deciding and intending.

Decision Intention Principle for Actions: If X decides to A, then X intends to A.

This view suggests a very close link between decisions and intentions. The holist accepts this link between decisions and intentions, but argues that just as one decides on scenarios, so one intends scenarios. According to Broad Holism, the holistic nature of decisions carries over, or transfers, to intentions. So the holist supports the following:

Decision Intention Principle for Scenarios: If X decides in favor of a scenario, then X intends that scenario.

I have been assuming so far that only that actions are intended. Now we have a proposal which assumes scenarios are intended. But perhaps this is not a problem. After all, scenarios just are collections of actions. So perhaps

we can understand 'X intends S' where S is a scenario consisting of X's A-ing and B-ing as 'X intends to A and B'.

The Holistic Decision Principle and the Decision Intention Principle for Scenarios (henceforth 'the Decision Intention Principle') together present the beginning of a threat to the DDE. The DDE proponent asserts that the strategic bomber does not intend to bring about the civilian deaths. The Holistic Decision Principle and the Decision Intention Principle together tell us that even the strategic bomber intends a scenario which includes bringing about civilian deaths. This of course leaves open whether the bomber intends to bring about the deaths themselves (as opposed to simply intending the scenario which includes the deaths). These two principles are compatible with the claim that X intends a scenario but does not intend any one part of the scenario and with the claim that she intends some parts but not others, thus they are compatible with the claim that the terror bomber intends to bring about civilian deaths and the strategic bomber does not intend to bring about the deaths. The DDE proponent could agree to the Holistic Decision Principle and the Decision Intention Principle. She could allow that bringing about e be part of an intended scenario without itself being intended.³ For example, so far the holist, if right, has shown only that the strategic bomber intends a scenario which includes his bringing about civilian deaths, and this differs from his intending to bring about the deaths themselves. So the DDE proponent could claim that there is a moral distinction between intending to bring about e and intending a scenario which has as a part your bringing about e.

³Roderick Chisholm appears to accept this view, or one much like it. Chisholm (1970).

Those who hold the position I am calling "Broad Holism" reject this option. They accept the Holistic Decision Principle and the Decision Intention Principle but object to the proposal above. Rather they endorse the following:

Intention Division Principle: If X intends scenario S and S includes X's bringing about e, then X intends to bring about e.

According to the Intention Division Principle, if the bomber intends a scenario which includes his bringing about the civilian deaths, then he intends to bring about the deaths. Thus the three principles above (the Holistic Decision Principle, the Decision Intention Principle, and the Intention Division Principle), entail that both bombers intend to bring about the civilian deaths.

Broad Holism then is the view which favors these three principles:⁴

Broad Holistic Decision Principle: If X foresees that A-ing will result in e and decides to A, then she decides on a scenario which includes A-ing *and* bringing about e.

Decision Intention Principle for Scenarios: If X decides on a scenario, then X intends that scenario.

Intention Division Principle: If X intends scenario S and S includes X's bringing about e, then X intends to bring about e.

⁴ See also Michael Bratman's discussion of the problem of the package deal in chapter 10 of Intentions, Plans, and Practical Reason. He presents the holistic view as an explicitly normative one; otherwise my presentation differs from his in fairly minor ways.

Broad Holism is incompatible with the Ordinary View. According to Broad Holism, an agent intends to bring about each effect of her voluntary action which she foresees. Thus, there is no difference between those effects an agent expects to bring about and those she intends to bring about. But that is to say that the Ordinary View is false. Since DDE depends upon the Ordinary View, Broad Holism presents a challenge to DDE.

In the next section, I will turn to a couple of cases in which Broad Holism yields a very counterintuitive result.

3. Two Objections to Broad Holism

The first case involves trivial effects. Ralph deliberates over whether to move from New York to San Francisco. He foresees that if he moves to San Francisco, he will need to find a new tailor. (He has no desire to bring his tailor with him.) Ralph's relationship with his tailor is quite mundane and his searching for a new one is an utterly trivial aspect of his moving to the West Coast. Nevertheless, it is a consequence he foresees. Thus, according to Broad Holism, if Ralph decides on the scenario in which he moves to San Francisco, Ralph intends to bring it about that he need a new tailor. But this is counterintuitive. Ralph is concerned with where he will live in San Francisco, how relocating will affect his friendships, what his new job will be like, how he will sell his condo in New York, and so on. The fact about his tailor plays no role in his deliberations; it does not influence his decision in any way. It is counterintuitive to think that he intends to bring about a need for a new tailor.

If this example is not entirely convincing, consider some other examples. If Broad Holism is true, then we are faced with the counterintuitive result that the agent who sets out on a walk across a field,

foreseeing that she will bend blades of grass, intends to do so. And the person who goes for a bike ride, foreseeing that by doing so he will burn calories, intends to do so.⁵ But surely this need not be the case.

A second sort of case poses even more serious problems for the Broad Holist. Consider an administrator who is responsible for hiring someone to fill a job opening. She believes it is her duty to hire the applicant with the best qualifications. She learns that one of her neighbors is applying for the job and foresees that her hiring decision will affect her relationship with her neighbor in the following way: if she gives the neighbor the job, their relationship will improve. If she gives the job to someone else, their relationship will suffer. The administrator does not want her relationship with her neighbor to suffer but believes that she must fulfill her duty to hire the best candidate. Wanting to be fair, she decides to screen-off (or exclude) all considerations about how the hiring decision will affect the relationship from her deliberations.⁶

If the administrator decides on the scenario in which she hires her neighbor, then according to Broad Holism, in spite of her 'screening off' efforts--she still intends to bring about an improved relationship with the neighbor. Here then is another counterintuitive result of Broad Holism. Intuitively it seems she does not intend to bring this result about. If someone comes to her and accuses her of being biased in her hiring choice, she might honestly respond: "My intention was to hire the best person for the job, not to benefit myself or my neighbor in any way." This case suggests that people do not intend parts of a scenario which they screen-off from their deliberations.

These two cases are meant to undercut the intuitive plausibility of Broad Holism. But a more plausible version of Holism can be formulated by

⁵Michael Bratman "Two Faces of Intention", (1984), p. 395.

⁶Alison McIntyre discusses such case (MS, 1987). The expression "screening off" is hers.

rejecting the Broad Holistic Decision Principle and replacing it with something else. So I will no longer consider the Broad Holistic Decision Principle, but will now turn to the second form of holism.

4. Deliberative Holism

The view I call "Deliberative Holism" differs in only one way from Broad Holism. The Deliberative Holist accepts the Intention Division Principle and the Decision Intention Principle, but replaces the Broad Holistic Decision Principle with what I will call the "Deliberative Holistic Decision Principle". Thus Deliberative Holism is the view which accepts these principles:

Deliberative Holistic Decision Principle: If X believes that e is an upshot of A-ing and *assigns weight to that fact in her deliberations*, then she decides to A if and only if she decides on a scenario which includes A-ing *and* bringing about e.

Decision Intention Principle for Scenarios: If X decides on a scenario, then X intends that scenario.

Intention Division Principle: If X intends scenario S and S includes X's bringing about e, then X intends to bring about e.

Let's look more closely at the Deliberative Holistic Decision Principle. What does it mean for an agent to *assign weight* to the fact that an effect will be brought about? Well, let's reconsider the hiring case discussed above. The administrator is aware of the following: (i) if she hires her neighbor, she

herself will have to do more of the computer work; and (ii) if she hires her neighbor, her relationship with her neighbor will improve. She allows (i) to bear on her decision about what to do. Since she does not want to do more computer work herself, this fact counts as a cost of hiring her neighbor; that is, she assigns it negative value. In contrast, she does not allow (ii) to enter her deliberations about what to do; she does not count it as a cost or benefit of hiring her neighbor. So assigning weight to the fact that a foreseen consequence will occur if one A's involves seeing this fact as a cost or benefit of A-ing and deciding to allow the cost or benefit to bear on the decision of whether to A.⁷

So far, Deliberative Holism may seem like an *ad hoc* response to the objections raised against Broad Holism. But I think the view is not *ad hoc*. Rather, the view emphasizes the link between intentions and deliberation. It would not be surprising if what one brings about intentionally is closely tied to what one brings about deliberately (in fact, it might be surprising if there were a difference between these), nor would it be surprising if what one brings about deliberately is closely tied to what enters into one's deliberations. The Deliberative Holist says we intend to bring about those effects to which we assign weight in our deliberations. The strategic bomber gives the deaths a negative value; he sees them as weighing against dropping the bomb. Screened off effects and trivial effects of acting play no role in an agent's deliberations and thus are not brought about deliberately.

⁷ McIntyre writes that the consequences the agent intends are the ones which are for the agent a reason for or against going ahead with the plan. She writes: "A reason against a plan is not merely some cost or consequence which would tend to decrease the plan's attractiveness. It must be a consideration which counts to the agent as a reason against adopting or going ahead with the plan in question. Some costs are too small to be seen as a reason against a plan..." (p. 28). I will assume that costs which are too small to count as a reason simply are not given weight in the deliberative process. I will treat costs which McIntyre calls too small as trivial consequences (ie ones which are assigned no weight by the agent).

Unlike Broad Holism, Deliberative Holism is compatible with the Ordinary View (p. 5). According to the Deliberative Holist, an agent may foresee but not intend to bring about trivial consequences and screened off consequences. If Deliberative Holism does not pose a challenge to the Ordinary View, why are we concerned with it? In the following section, I discuss why a proponent of DDE would be unhappy with Deliberative Holism.

5. Deliberative Holism and DDE

While Deliberative Holism does allow for a distinction between what an agent intends to bring about and what she merely foresees she will bring about, it does not allow for the distinction needed by DDE. For example, the DDE proponent typically assumes that the pilot in STRATEGIC BOMBER foresees he will cause deaths without intending to do so. But according to Deliberative Holism, since the strategic bomber assigns weight to the deaths, it will turn out that he (like the terror bomber) intends to bring the deaths about. So if Deliberative Holism is true, DDE will no longer draw a moral distinction between STRATEGIC BOMBER and TERROR BOMBER.

There will be other results which are problematic for DDE. Imagine a third bomber, the callous bomber. The callous bomber is in much the same situation as the strategic bomber. He wants to weaken the enemy by causing the destruction of the factory and he foresees that by bombing he will cause civilian deaths as well as the destruction of the factory. The strategic and callous bomber differ in this way: the strategic bomber sees the civilian deaths as a cost of bombing while the callous bomber does not. The callous bomber assigns no value to the deaths; he is completely indifferent to this expected

effect of bombing. According to Deliberative Holism, the strategic bomber intends to bring about civilian deaths while the callous bomber does not.

Furthermore, according to DDE, the strategic bomber, who worries about the deaths, may not drop his bomb, but the callous bomber, who does not care about the deaths, may permissibly drop his bomb. This surely cannot be right.

I offer this case not as an argument against Deliberative Holism, since the Deliberative Holist may simply say of it: "So much the worse for the Doctrine of Double Effect." Rather the case is meant to show that Deliberative Holism, if true, diminishes the plausibility of DDE and thus to show why a proponent of DDE would be motivated to argue against Deliberative Holism.

6. Objection to Deliberative Holism

Although Deliberative Holism does not yield counterintuitive results where trivial effects and screened off effects are concerned, it will yield counterintuitive results in other cases. To see the problem, we should look at a case which does not bring in DDE. Consider the case of Frieda. Frieda is a surgeon. While she is on duty in the emergency room, a very ill patient is brought in. If Frieda does nothing, the patient will die. If Frieda operates, there will be the following two results: the patient will live and the patient will be paralyzed. For Frieda, bringing about the paralysis weighs strongly against operating. Still, saving the patient's life weighs more strongly on operating, so she decides to operate.

According to the Deliberative Holist, Frieda intends to bring about the patient's paralysis. After all, Frieda assigns negative value to this foreseen effect of operating; it is a reason against operating. But surely she does not intend to bring about this effect.

The Deliberative Holist would presumably acknowledge the counterintuitiveness of the claim that Frieda intends to bring about the patient's paralysis. It is not at all surprising that such cases can be generated given that the Deliberative Holist holds that an agent may intend to bring about even those effects she sees as costs of acting. The Deliberative Holist could stress that although bringing about the paralysis is a cost of acting for Frieda rather than something she aims at, it is a cost which Frieda acknowledges and takes responsibility for bringing about. Alison McIntyre defends Deliberative Holism. She claims that reasons against acting are "costs of acting which the agent acknowledges and so takes a certain kind of responsibility for, in acting."⁸

My view is that an account of intention which entails that Frieda intends to bring about paralysis should be rejected. An account of intention which strays too far from our ordinary understanding will not help the holist; it will simply leave room for the DDE proponent to argue that the morally important distinction is between intending in the ordinary sense rather in this new, broader sense and merely foreseeing harmful consequences. The question then is which of the three principles constitutive of Deliberative Holism should be rejected.

7. Rejecting the Holistic Challenge

A proponent of the Ordinary View will of course want to reject at least one of the three principles. To review, those principles are:

⁸McIntyre, p. 28.

Deliberative Holistic Decision Principle: If X believes that e is an upshot of A-ing and *assigns weight to that fact in her deliberations*, then she decides to A if and only if she decides on a scenario which includes A-ing *and* bringing about e.

Decision Intention Principle for Scenarios: If X decides on a scenario, then X intends that scenario.

Intention Division Principle: If X intends scenario S and S includes X's bringing about e, then X intends to bring about e.

I'll begin by discussing the Intention Division Principle. Why think that if X intends a scenario, she intends each part of the scenario as well? If an agent intends scenario S, then either she intends no parts of S, she intends some parts of S and not others, or she intends all parts of S. No one wants to claim that she intends no parts of the scenario. This would entail that an agent does not intend to bring about e even where doing so is her end or a means to her end.

But what about the view that she intends some parts and not others? If we grant the holistic nature of decisions, then to preserve commonsense distinctions, we might reject the Intention Division Principle and maintain that only some parts of a scenario are intended. I want to consider a couple of the most obvious proposals for making this distinction between parts of the scenario which are intended and parts which are not.

First Proposal: *If X has decided on a scenario which includes bringing about e, then X intends to bring about e if and only if she wants to bring about e.* This proposal locates the difference between intending and merely

foreseeing in desire. I think though that wanting to bring about e is neither necessary, nor sufficient for intending to bring about e. Let's consider necessary condition first.

We can imagine that the terror bomber does not want to bring about any deaths. Perhaps he hates harming others and if you ask him later, 'Did you want to cause those deaths?', he will say 'No, that was the hardest thing I have ever had to do'. Even in the absence of a desire to cause the civilian deaths, the terror bomber intends to cause civilian deaths.

The DDE proponent might try to respond by claiming that there are narrow and broad senses of 'want' or 'desire' (call these 'want_n' or 'desire_n' and 'want_b' or 'desire_b'). While the terror bomber does not want_n to bring about deaths, he nevertheless wants_b to bring them about; after all he believes the deaths will bring about his end.

But the holist has a response; the holist will point out that once we admit some broader notion of desire, then the strategic bomber may also be said to want_b to bring about deaths. The strategic bomber wants_b to bring about deaths in so far as he takes doing so to be connected to his end. This response is at first unconvincing. After all, it may be that the strategic bomber attaches negative value to bringing about the deaths, but that he attaches enough positive value to other elements of the scenario that the scenario as a whole is preferable to alternatives. But notice that the same can be said of the terror bomber. Perhaps he too attaches negative value to bringing about the deaths but prefers the scenario which includes bringing about the deaths to other options.

So far I have been arguing that desire fails to provide a necessary condition for a part of a scenario being intended. I think desire also fails to provide a sufficient condition for a foreseen effect to be intended. There seem

to be cases where although the agent does desire a foreseen side effect, the nonholist would want to say that the foreseen effect is unintended. The following may be such a case:

Frederico loves to sing at night. His singing greatly irks his neighbor and he knows it. In fact, he is quite pleased with this upshot of his nocturnal practicing. Still, if we ask him why he sings at night, he will tell us that his voice is most relaxed and at its peak in the night. If we ask him whether he intends to annoy his neighbor, he will adamantly deny that this is his intention. He looks upon this effect as an unintended benefit of singing at night.

One might attempt to respond to this example by claiming that Frederico intends to practice his singing and to annoy his neighbor. Certainly, he might intend to do both; he might sing a little louder if he's not sure he is keeping his neighbor awake, and so on. But we need not assume he has the intention to do both.

We have been looking for a way to distinguish parts of the scenario the agent intends from parts he does not intend. I have been arguing that an agent's desiring a part of a scenario does not provide a sufficient condition for his intending that part. I have also been arguing that appeals to desire as a necessary condition for intending don't work. I want to turn now to a second proposal.

Second Proposal: *If X has decided on a scenario which includes bringing about e (as a result of A-ing), then she intends to bring about e if and only if bringing about e is a means or end (and not a side-effect of A-ing). We*

often talk about side effects as unintended and means and ends as intended. The DDE proponent might point out that although the strategic bomber intends a scenario which includes bringing about the civilian deaths, he does not intend to bring about the deaths themselves since he views bringing about the deaths as a side effect of his bombing and not as a means to his end.

But this proposal will not do. Whether or not an action is a means to a given end is irrelevant to whether the agent intends that action as a means to the end. For example, assume the terror bomber is mistaken in believing that civilian deaths will lead to a drop in enemy morale. In fact, civilian deaths will not causally contribute to his end of weakening the enemy. This does not entail that he does not intend to cause the deaths.

Still, it seems the original proposal is easily modified. What we want is the following: *If X has decided on a scenario which includes bringing about e (as a result of A-ing), X intends to bring about e if and only if she believes that bringing about e is a means or end (and not a side-effect of A-ing) .*

Before we can evaluate the modified proposal, we need to have an understanding of what is meant by "side effect". If by "side effect", we just mean an effect which the agent does not intend to bring about then appeal to side effects versus means and ends will be of no help; it will be circular. What is needed is an account of side effects versus means which does not itself rely upon the notion 'intending'. Giving such an account may be surprisingly difficult. For example, one might say that the deaths are a side effect for the strategic bomber because he does not really need to bring about deaths to achieve his end. One might point to a counterfactual to make the point: if the strategic bomber believed that as a result of his raid, the armory was destroyed but amazingly, the civilians were unharmed, he would not circle back for another attack. But Jonathan Bennett has pointed out a problem with such

reasoning. Bennett points out that by the same sort of reasoning one can show that the terror bomber does not really intend to kill civilians. After all, if the terror bomber believed the civilians would miraculously recover after the war, he would be no less motivated to bomb. So perhaps he too does not intend to cause deaths. (This is a crude version of Bennett's argument. I discuss Bennett's objection at much greater length in Chapters 2 and 4.)

We have been looking for a way to distinguish parts of the scenario which are intended from parts which are not (in order to block the holist's claim that all parts are intended) in order to reject the Intention Division Principle. My conclusion here is not that there is no way to do so, but that there is no obvious and unproblematic way to make the distinction. Even so, given this weak conclusion and the fact that no one wants to claim that no parts of the scenario are intended (surely agents intend to bring about their ends), we need to take the Intention Division Principle seriously.

8. Bratman's Response to Deliberative Holism

In Intentions, Plans, and Practical Reasoning, Michael Bratman considers a view very similar to Deliberative Holism.⁹ He endorses a principle similar to my Intention Division Principle. Also, Bratman endorses something fairly close to the Deliberative Holistic Decision Principle. He accepts the following:

If I know my A-ing will result in e and I seriously consider that fact in my deliberations about whether to A and still go on to conclude on A, then if I am rational my reasoning will have

⁹In Chapter 10 of Bratman (1987), Bratman discusses what he calls 'the problem of the package deal'. He represents the problem as issuing from the acceptance of four principles, rather than three and the principles he discusses are all normative.

issued in a conclusion on an overall scenario that includes *both* my A-ing *and* my bringing about e.¹⁰

This principle differs in certain ways from my Deliberative Holistic Decision Principle, most notably in that I have not made any qualifications for whether or not the agent is rational. Still, since Bratman accepts this principle, he clearly accepts some sort of holism about decisions.

Bratman rejects holism about intentions and does so by driving a wedge between intentions and these conclusions (which he calls 'practical conclusions' and I call 'decisions'). He accepts that practical conclusions are holistic in nature but rejects the view that this holism transfers to intentions. So if faced with the triad of principles above, he would reject the Decision Intention Principle.

Why think that decisions are holistic and intentions are not? According to Bratman, practical conclusions are subject to different sorts of pressures than are intentions. These pressures come from concerns about rationality. He writes that practical conclusions are holistic because they are subject to standards of 'clearheadedness' or intellectual honesty. For example, when Frieda deliberates about what to do, she cannot honestly and rationally ignore the fact that operating on the patient will cause paralysis. If Frieda is rational, her conclusion will not merely be to operate but will be in favor of the whole package—operating, saving the patient, and bringing about the patient's paralysis.

This is a normative view. It says it would be irrational not to take into account all foreseen consequences of acting. My interest here is in how we

¹⁰ Bratman (1987), 145

actually make decisions, but it's highly likely that these come to the same thing (in at least most instances).

Deciding on a scenario involves summing up all the costs and benefits. These costs and benefits, according to Bratman, are 'backward-looking' pressures while intentions are subject to 'forward-looking' pressure. An agent's deliberations (should or on the view I am considering, do) reflect what she believes to be the facts--thus, Frieda cannot rationally ignore the paralysis when deliberating. But the connection between intentions and prior deliberation is, according to Bratman anyway, a looser connection. Intentions are individuated by their role in further reasoning and action--they are future-directed. They are not subject to the same standards of 'deliberative honesty'.

I find it difficult to grasp Bratman's meaning here. A practical conclusion is supposed to differ from an intention in being backward, rather than forward, looking. But practical conclusions are conclusions to do such-and-such. A conclusion to A or to bring about e is surely forward-looking (or at least present-directed) rather than backward-looking. So I am suspicious of the attempt to drive a wedge between these conclusions and intentions.

Also, I find his point about deliberative honesty unconvincing. Surely Frieda can decide simply to operate while meeting standards of rationality and deliberative honesty. She does not ignore her belief that operating will bring about paralysis. She takes this consequence seriously--she sees it as weighing against operating. This is sufficient to make her reasoning clearheaded. She does not need to include the consequence in the content of her conclusion, or choice. We can distinguish an agent's reasons for reaching a conclusion from the content of that conclusion.

I find Bratman's reasons for rejecting the Decision Intention Principle unconvincing. In the following section, I too will argue for the rejection of this principle, but on grounds different from those offered by Bratman.

9. An Alternative Rejection of Holism

The Decision Intention Principle asserts that if an agent decides on a scenario, that agent intends the scenario. My view is that when one decides on something, one is reaching a theoretical conclusion; one is making an evaluative judgment. Notice that we decide on all sorts of things, not just collections of actions. For example, we decide on blue shirts over red ones, New England Boiled Dinners over Burritos, and so on. There may be reason to treat such decisions in a uniform way. Assuming that deciding on S entails intending S rules out a uniform treatment. After all, if you decide on the New England Boiled Dinner, you do not intend the dinner. Deciding on the dinner may give rise to some other intention, like an intention to buy the dinner, but that's another matter.

My suggestion is that we distinguish *deciding to* from *deciding on*. *Deciding on* may correspond to reaching a theoretical conclusion while *deciding to* may correspond to reaching a practical conclusion. My view is that if one decides *to* A, then one intends to A. However if one decides *on* something, one does not intend that thing. This allows us to treat deciding on something in a uniform way, whether what is decided on is a collection of actions or a shirt. In all cases, such a decision is an evaluative judgment. Such judgments often inform or give rise to intentions, but not always. The Decision Intention Principle sounds plausible because scenarios include only actions and actions are the sorts of things agents can intend. But I think we

should see that to decide on a scenarios is still just to make a preference-ranking; it's not yet to decide to *do* anything.

There is a further consideration which may support my view and suggest that this line of response is one Bratman might have taken as well. Bratman asserts that practical conclusions are conclusions which we decide upon whereas theoretical conclusions are ones we discover.¹¹ When one ranks preferences as above, one is engaging in what Bratman would call 'desire-belief weighing'. But notice what this involves. It involves taking into account one's desires, one's beliefs (for example, of the likelihood that a given action will have a certain outcome) and then doing some calculations. We do not decide what to desire anymore than we decide whether to be in pain. What we desire, like what we feel, is something we discover. This may sound odd, but I think the oddness is attributable to the fact that discovery typically involves the risk of error and these cases are not ones where we go wrong. Nevertheless, we do not decide what to desire. Likewise, we do not decide what to believe. So the process of desire-belief weighing really falls into the category of theoretical reasoning, if one understands theoretical conclusions as conclusions we discover. I think all this supports my position that to decide on a scenario is not yet to intend to do anything.¹²

Let me summarize the two responses to the holistic challenge which I have been contrasting. Bratman rejects the holistic challenge by driving a wedge between practical conclusions and intentions. He defends this position

¹¹He also says that practical conclusions are conclusions concerning what to do. So perhaps he takes these to coincide. I think reasoning about what to do includes preferences rankings, beliefs, and other sorts of theoretical conclusions.

¹²The view I am asserting here does I think face a serious problem. If desire-belief weighing is not deciding then what is? This view may make cases of ordinary decision-making too mysterious. (The flip side of the problem is this account may also too easily explain weakness of will.)

with the suggestion that practical conclusions are backward-looking and intentions are forward-looking. I think practical conclusions (or decisions to act) are forward-looking. I don't want to drive a wedge between these and intentions. I do agree with Bratman that we can reject the Decision-Intention Principle, but because I think that deciding on a scenarios is not reaching a practical conclusions. These conclusions are not yet conclusions to do anything. Thus while I think the following is true "If X decides to A, then X intends to A", I think the Decision Intention Principle (for Scenarios) is false.

In this chapter I have considered and rejected two versions of holism, each of which threatened to undermine the distinction between intending and foreseeing which is assumed by the Doctrine of Double Effect. In part, I have argued that deciding on a scenario may be merely evaluative and need not usher in an intention. In the following chapter I turn to a second challenge to the view that the terror bomber intends to bring about civilian deaths and the strategic bomber does not.

Chapter Two: Bennett's Challenge

1. Introduction

In the previous chapter, we saw that some of the most natural ways to distinguish what is intended from what is merely foreseen fail. Recall that appeals to desire to mark the difference did not succeed and I suggested that there was a serious problem with appeals to the notions 'means' and 'side effects' as well. In this chapter, I will be turning to that problem. But first, I will take a brief look at two different types of accounts of intention.

2. Reductive Accounts of Intention

Elizabeth Anscombe, Donald Davidson (in his early work) and others advocate what might be called a reductive view of intentions.¹ Advocates of this sort of view tend to focus on intentional action. What makes an action intentional is its relations to certain beliefs and desires of the agent. For example, what makes it the case that Max intentionally cracks the egg will be facts about the relations his egg-cracking bears to his beliefs and desires. For example, he wants an omelette and believes that cracking the egg is a way of making an omelette. Otherwise put, his desire for an omelette and his belief that cracking the egg will help to produce an omelette give him a reason to

¹The different reductive views vary from each other of course. Perhaps most notably they differ in whether they assert a causal relation between the action and the relevant desires and beliefs. I ignore all these differences since they are not pertinent here.

crack the egg. Such views tend to deemphasize future directed intentions (intending). And it is really future directed intentions which are relevant to the Ordinary View. But some contemporary advocates of the reductive view extend the account given of intentional action to future directed intentions. It is common on such views to see intending as a complex of believing and desiring (for example, Sally intends to go skiing if and only if she wants to go skiing more than she wants to do anything incompatible with skiing and she believes she is more likely to ski than not to ski, or as a special kind of desire).²

3. A (non-reductive) Planning Account of Intention

In Intentions, Plans, and Practical Reason, Michael Bratman presents a detailed account of intentions which differs in important ways from reductive views. While proponents of reductive views typically reach their conclusions by considering cases of acting intentionally, Michael Bratman arrives at his account by starting with cases of intending to do something in the future (for example, intending to drive to Boston tomorrow). The difference is important. Bratman takes these future-directed cases of intention to be paradigm cases. In such cases, there is no action. For example, I may intend to drive to Boston tomorrow, yet currently, I am sitting at my computer typing. My intention to drive to Boston tomorrow is not manifested in my current behavior, nor does it explain anything about my current behavior. (Of course it might explain why I am typing this sentence now and not tomorrow, but it need not.) These future-directed cases of intending are problematic for reductive accounts like Anscombe's and

²This is a highly simplified version of Robert Audi's view. See Audi (1970).

Davidson's. Even if acting intentionally is simply performing an action which is related in a certain way to one's beliefs and desires, where does this leave merely intending to A? (One cannot give the same explanation since the action is missing.) Bratman argues that attempts to extend accounts of intentional action future directed intentions fail. Rather, he thinks, we need to recognize intentions as a distinctive state of mind, on a par with beliefs and desires.³ I will not attempt to summarize or defend Bratman's arguments for this conclusion here. Rather, I am interested in the resulting account of intentions which he advocates. If intentions are to be taken seriously, on a par with beliefs and desires, what are they?⁴ I devote more attention to Bratman's planning view than to the reductive accounts. I do so because, as will be seen shortly, Bennett relies upon an Anscombe-like view of intentions to make his objection. It will be important to ask whether Bennett's objection holds on other accounts of intention.

Bratman thinks that to understand intentions, we need to see them as elements in plans. Intentions are what allow both interpersonal coordination and intrapersonal coordination. For example, one person is able to coordinate meeting another at the airport partly by forming an intention to be at the airport at a certain time. This intention, once formed, has a certain sort of stability. If a person intends to be at the airport Thursday morning, she will not continue deliberating about what to do Thursday morning; rather, she will see the issue as settled. Likewise, she will not make plans which she

³In later work, Donald Davidson arrives at a similar conclusion by considering future-directed cases of intending. See 'Intending'.

⁴One objection to this approach might be that attitudes can be divided exhaustively according to what Anscombe called their 'direction of fit'. Desires and desire-like states have one direction of fit, beliefs have the other. Intentions apparently have the same direction of fit as desires, so why not see them as a special sort of desire? Because intentions differ in an important way from desires. There is nothing criticizable about having desires one knows to conflict, however there is something criticizable about having intentions one knows to conflict.

thinks conflict with her being at the airport on Thursday morning. Intentions are to be understood in terms of the relations they bear both to an agent's reasoning and to an agent's behavior. Bratman sees intentions as having what he calls a 'reason-centered dimension of commitment' and a 'volition-centered dimension of commitment'. So someone who has formed the intention to go to Boston tomorrow will see the issue as settled. She will not be disposed to continually reconsider whether or not she should go. This is the reason-centered aspect of commitment. Likewise, the intention is conduct-controlling. A regular desire is not conduct-controlling. Though I would like some chocolate right now this desire does not influence my behavior. An intention to finish typing this sentence, will on the other hand, influence my behavior. This is the volition-centered aspect of intention.

Bratman suggests that intentions characteristically play the following three roles, the first two partly constituting the reason-centered dimension of intentions, the third partly constituting the volition-centered dimension of intentions.

An intention to do alpha will characteristically play these roles:

(R1) drive means-end reasoning regarding how to do alpha,

(R2) act as a background filter, screening out options the agent believes to be incompatible with doing alpha, and

(R3) lead to endeavoring to do alpha.

Again, Bratman does not suggest that (R1)-(R3) are necessary or sufficient conditions for intending. Rather, he says they are roles *characteristic* of intentions. Let's take a closer look at each of the conditions above.

According to (R1), once one has formed an intention to A, one then needs to reason about how to do it. For example, once I form an intention to go to California, I will then turn to the question: how am I going to do it? Am I going to fly and if so, from what airport. Upon forming the intention to leave from a certain airport, I will then turn to the question of how I'm going to get to the airport. But what about the following sort of case: I form an intention to go to work this morning. Assuming I take the same route to work every day, I will not be disposed to reason further about how to get to work. Does this show I lack the intention to go to work? No. Notice that if my usual route is blocked by construction, then I will reason about the best alternative route. Once one forms an intention, one is disposed to engage in means-end reasoning. Whether one actually does engage in such reasoning will depend upon a number of factors about the particular situation, most notably, how complex carrying out the intention is and how many times one has done something similar before.

According to (R2), if I am intending to meet you at the airport on Thursday morning, then I will not also agree to have coffee with someone else, somewhere else, at the same time. I will not to make plans which I believe to be incompatible with my being at the airport Thursday morning at ten. Imagine that a work associate tells you she will have lunch with you the following day. Moments later you hear her making plans to have lunch with someone else (somewhere else), also for the next day. Assuming you have reason to believe that she is not planning to have two lunches tomorrow, you will doubt whether she really intends to have lunch with you. This is because if one intends to A, then one typically screens out options believed to be incompatible with doing A.

Endeavoring has to do with the link between intention and action. To endeavor is to adjust one's behavior in response to signs of success or failure. If I intend to meet you at the airport at ten a.m. and I see that I am running late, I will speed up. Likewise, if I see that I have taken a wrong turn and am no longer headed toward the airport, I will turn around so that I'm headed in the right direction. Notice that even very simple intentions involve endeavoring. Imagine that while reading a book, you form an intention to have a sip of the coffee which is sitting on your desk. You reach for the coffee while keeping your eyes on the book. If your hand does not immediately make contact with the cup of coffee, you will move your hand around trying to find the cup, and at some point, look up and allow visual cues to guide your hand in the direction of the cup. So even this very simple case of intending to have a sip of coffee involves endeavoring.

Now we come to the question of whether Bratman's view helps us to reject Bennett's challenge. Bratman believes that his account does distinguish the two bombers; that is, he believes his account supports the view that the terror bomber intends to cause civilians deaths and that the strategic bomber does not. Let's apply (R1)-(R3) to the two bombers and see how this is supposed to work.

The terror bomber will be disposed to reason further about how to bring about civilian deaths. Perhaps he will deliberate over the height from which he should drop the bomb in order to cause the most casualties. The strategic bomber, on the other hand, will not reason further about how to cause civilian deaths; perhaps he will be disposed to think about the best height from which to bomb in order to ensure the total destruction of the factory.

The terror bomber, but not the strategic bomber, will rule out options which are incompatible with causing civilian deaths. Bratman elaborates on this by imagining that after settling on his plan to terrorize by causing civilian deaths, the terror bomber considers ordering a certain troop movement. He sees that this troop movement, while having certain military advantages, would also alarm the enemy and cause them to evacuate the civilians he plans to kill. Since moving the troops in this way has a consequence which he sees is incompatible with his prior intention, he will not form a new intention to move his troops in this way.⁵ Bratman sees the prior intention as forming a screen of admissibility. An option which is incompatible with an already formed intention will not pass through the screen.⁶

Notice that an intention to A not only leads one, in the typical case, to refrain from forming intentions incompatible with A, it also leads one to refrain from acting in ways which are incompatible with A-ing.

Finally, let's look at endeavoring. If the terror bomber were to see his bomb did not in fact cause any civilian deaths, he would bomb again. If the strategic bomber were to see that his bomb did not cause civilian deaths (but did cause the destruction of the munitions plant), he would not bomb again. Likewise, if the civilians were to run away from the armory upon hearing the plane, the terror bomber would follow them, but the strategic bomber would not.

According to Bratman, the terror bomber reasons about how to kill the civilians, he screens out options believed to be incompatible with doing so,

⁵Bratman (1987), p. 141

⁶There are exceptions, new information can lead us to revise an intention we already have; I am putting aside such cases, but Bratman's view does accommodate intention revision.

and he endeavors to kill them. Not so the strategic bomber. Thus only the terror bomber intends to kill the civilians.

4. Counterfactual Test Question

Let's turn now to Bennett's argument. Bennett assumes Anscombe's account of intentions is the correct one. (I believe he has reductive views generally in mind here.) So Bennett suggests that the difference between the two bombers is that the terror bomber is in some way motivated by his expectation that the raid will produce deaths while the strategic bomber is not so motivated by the expectation of deaths. Bennett suggests the following: "If intentions are determined by which of the person's beliefs motivate his action, then we should be able to get at them by asking how the behavior would have differed if the beliefs had differed in given ways." He then offers the following, at least initially very plausible, counterfactual test question:

Counterfactual Test Question: If you had believed that there would be no civilian deaths, would you have been less likely to go through with the raid?⁷

The idea here is that a positive response to the Counterfactual Test Question indicates an intention to bring about the deaths whereas a negative response indicates the lack of such an intention. In considering how each bomber would respond, we need to know what else is supposed to vary along with the belief about civilian deaths. We are looking for an interpretation under which the bombers answers to the test question differ. Bennett suggests

⁷This is the test question suggested by Jonathan Bennett is part III of his Tanner Lectures (1981).

that there are three possible interpretations and that the bombers answers differ only on the third of these interpretations. He reasons as follows.

First Interpretation: Let's assume that the bombers in the counterfactual situation differ from those in the actual situation only in this way. In the counterfactual situation what varies is each bomber's belief that the raid will cause deaths and any other of their beliefs which is logically entailed by the belief that the raid causes no deaths. So under this interpretation since each bomber believes no deaths result from the bombing, each believes as well that no one is killed as a result of the bombing. In this way, we do not need to assume that the bombers in the counterfactual situation have inconsistent beliefs if they don't have inconsistent beliefs in the actual world.

Bennett claims that this way of interpreting the test question will not yield the desired results since on this interpretation both bombers answer the test question No. He suggests that the strategic bomber answers No to the test question since the strategic bomber retains his belief that the weapons plant is destroyed. The terror bomber answers No because he retains his belief that there will be terror as a result of the raid. So neither would be less motivated to bomb.

It's puzzling why Bennett thinks that the strategic bomber retains his belief that the weapons plant will be destroyed and that the terror bomber retains his belief that there will be terror as a result of the raid. For example, doesn't the strategic bomber think that the weapons will be destroyed only if the civilians die? If so, then his believing that there will be no deaths (taken with the conditional belief: the factory will be destroyed only if there are deaths) will entail that the factory is not destroyed. Perhaps Bennett simply does not consider such conditional beliefs. In the end, it won't matter.

If we do count such beliefs, it will turn out that under this first interpretation both bombers answer Yes. So either way this is not the desired interpretation. Bennett then considers a second interpretation.

Second Interpretation: Perhaps what varies in the counterfactual situation is not only each bomber's belief about whether civilian deaths result but also beliefs about anything the bombers take to be causally related to the deaths. In that case, the terror bomber believes not only that there will be no deaths, but also that there will be no terror, since he takes the terror to be causally related to the deaths. So under this interpretation, he answers the test question Yes.

What about the strategic bomber? He has the belief that the factory can't be destroyed without killing the civilians. So he will suppose not only that there will be no civilian deaths but also no factory destruction. In this case he too will answer the counterfactual test question Yes. Again, we do not have the desired interpretation.

Third Interpretation: Bennett points out that in order to get the test to give the desired result (that the strategic bomber answer No to the test question and the terror bomber Yes) we need to assume that what varies in the counterfactual situation is the belief about civilian deaths and anything that follows from this belief by what Bennett has called "a causally downstream inference". (Where x causes y and y causes z, x is causally upstream from y while z is causally downstream from x. Thus a change in the agent's belief about y would affect the agent's belief about z, but not about x.)

Let's look again at the counterfactual test question this time assuming that what varies with the belief that there will be no civilian deaths is anything the bomber takes to be related to the deaths in a causally downstream way. The terror bomber believes that the terror is caused by the

deaths. So in the case where he believes there will be no civilian deaths, he also believes there will be no terror. Thus he will answer Yes to the test question. The strategic bomber's believes there will be no deaths but since he does not take the destruction of the munitions plant to be causally downstream from the deaths, he retains his belief that the factory will be destroyed. Thus, he will answer No to the counterfactual question.

It looks like we have found the desired interpretation of the test question. Again, the test question was of this form:

"If you had believed that there would be no event of type G as a result of your A-ing, would you have been less likely to A?"

And we are now assuming that, given the third interpretation of the test question, we are able to pry apart the two bombers. The test question suggests that the terror bomber is motivated by the expectation of civilian deaths and the strategic bomber is not, thus only the terror bomber intends to bring about the deaths.

5. Bennett's Reductio

It looks like we have found the interpretation of the test question which yields the desired results. But Bennett argues that in fact, even this interpretation, the one that looked all right, won't do, even it won't distinguish the bombers. We assumed that the terror bomber would answer the test question Yes (he would be less motivated to bomb). Bennett's strategy is to show that we made this assumption too quickly. He begins by asking us to consider the following "innocent" example:

A political leader takes action against a trade union, intending to bring about a month-long state of disintegration in which the various locals break off from the parent body and severally fall into further disunity. This is his intended means to the end of the union's being unable to call a strike during December. He is rightly sure that if the union falls apart for that long it will never be reconstituted, but all he cares about or intends is the one-month dissolution: if he were sure that the union would recover during January and flourish for many years, that would not reduce his motivation for moving against it.⁸

Bennett claims that the politician has 'killed the union' and done so knowingly, but that all the politician intends to bring about is a one-month dissolution. He argues:

- (1) If feature F is a conjunction of features G and H, you don't intend to produce an F unless you intend to produce a G and intend to produce an H.
 - (2) Killing the union at the beginning of December involves the union's being inoperative for all of December and the union's being inoperative from January first on.
 - (3) The politician does not intend to produce an event with the feature *union inoperative in January*.
- (C) Therefore, the politician does not intend to kill the union.

The same line of reasoning can be applied to the case of the terror bomber, with the analogous conclusion that the terror bomber does not intend to kill the civilians. That argument runs like this:

⁸Bennett (1981), 110

- (1) If feature F is a conjunction of features G and H, you don't intend to produce an F unless you intend to produce a G and intend to produce an H.
- (2') Killing the civilians involves making their bodies inoperative for a while and making their bodies inoperative thereafter.
- (3') The terror bomber does not intend to produce an event with the feature *civilians inoperative after the war*.
- (C') Therefore, the terror bomber does not intend to kill the civilians.

This argument is part of a reductio. Bennett is arguing that "the concept of what is 'intended as a means' cannot be given a firm, clear, theoretic grounding which implies what we think true and not what we think false regarding what people intend."⁹ His reductio starts by assuming such an account of the notion 'intend as a means' can be given. He says there is only one plausible account of intentions, according to which intentions are given by which beliefs motivate an agent. (Such an account is what supports (1) above). But when we apply the distinction between what is intended and what is merely foreseen, so understood, to given cases, we get absurd results. For example, we get the result that the terror bomber does not intend to bring about deaths. But this is absurd and since there is no alternative account of intentions available, we need to reject the assumption that the 'intend/foresee' distinction can be given a firm grounding; and thus, we should not rest moral weight on this distinction.

Let's return now to (1)-(C') above. In reconstructing this line of argument, I have stuck very closely to his words. For the conclusion (C') to follow from the premises, (2') has to mean that *being a killing* is a feature of some events, as are *inoperativeness for a while* and *inoperativeness*

⁹Bennett (1981), 113

thereafter. Furthermore, we need to assume that *being a killing* is a conjunction of the other two features. Some events are killings but it's not clear what it is for an event to have the feature *inoperativeness-for-a-while*. Perhaps Bennett has something like the following in mind. Some events bring about permanent changes. Perhaps Bennett's view is that *bringing about permanent inoperativeness which begins at t1* is a conjunction of *bringing about inoperativeness from t1 to t2* and *bringing about inoperativeness from t2 onward* (where t2 follows t1). And, if one does not intend to do something which brings about permanent inoperativeness, one does not intend to kill.

(1), which I will henceforth call 'the Disjunction Principle', is doing much of the work in the reductio and you might think there are obvious counterexamples to it. Take a simple case: Max intends to remain a bachelor. Won't Max intend to remain a bachelor simply by intending to remain single? Does he additionally need to intend to remain male? However, this case really isn't a counterexample to the Disjunction Principle. The example makes a claim about whether Max intends to G (to remain a bachelor), but the principle is not about whether one intends to G ; rather it is about whether one intends to bring about an effect of kind G. The difference turns out to be important. The latter formulation is more restrictive. It will only attribute to an agent an intention to bring about a G if the agent believes a G to be an upshot of one of her voluntary actions. Thus agents will not turn out to intend things over which they lack the requisite sort of control. Return to the example of Max. The Disjunction Principle would apply to a voluntary action which Max takes to have as upshots his remaining single AND his remaining male. One concocts a case in which Max is deliberating about whether to

perform some such act, but I suspect that once the relevant story were told, the case would not provide a counterexample to the Disjunction Principle.

Furthermore, there is reason to embrace the Disjunction Principle. After all, the strategic bomber knowingly produces both the factory destruction and the killing of civilians. We want to maintain that the strategic bomber intends to produce the factory destruction but does not intend to kill the civilians. The Disjunction Principle keeps us from attributing intentions which are too coarse (eg: the strategic bomber intends to produce deaths and destruction). Either we need to embrace the Disjunction Principle or to endorse some other account of intentions which allows that the strategic bomber does not intend to produce the deaths.

Bennett supports (3') by appeal to the counterfactual test question. The bomber would not have been less motivated to carry through with his raid had he believed the bombers would not be inoperative after the war, so he does not intend to cause their inoperativeness after the war.

From these premises, it follows that the bomber does not intend to cause the civilian deaths, that is, he does not intend to kill the civilians. Bennett's conclusion here is not that we should delimit the scope of intentions more narrowly than we are used to doing. Remember that the terror bomber case is supposed to provide a paradigm example of intending to kill as a means. Rather Bennett's conclusion is that on the only plausible account he can find of what it is to intend as a means, we get the absurd conclusion that the terror bomber does not intend to kill as a means, thus in the absence of a satisfactory account of the notion 'intend as a means', the notion should not bear moral weight.

6. Initial Responses to Bennett

Something can be said to soften the blow of the reductio. Bennett suggests that all the terror bomber intends is to bring about the civilians temporary state of inoperativeness. But this is misleading. After all, we ask, "how does the terror bomber plan to bring *this* result about?" Now Bennett might answer simply "by dropping his bomb on them" but if we push the point, it seems the bomber will tell us that dropping a bomb will cause the civilians to be blown apart and being blown apart will cause the inoperativeness (or perhaps it just is the inoperativeness). So the terror bomber intends not only to cause inoperativeness for a while, but also intends to cause the civilians to be blown apart. This may not soften the blow of the reductio much, but I do think it helps some. Still, this response leaves us with the troubling result that the terror bomber intends to cause civilians to be blown apart, but does not intend to cause their deaths.

Many people's response to the claim that the terror bomber intends to blow civilians to bits but not to kill them is that there is something illegitimate about driving a wedge between these. As Foot said, some effects are simply "too close" for the application of the Doctrine of Double Effect.¹⁰ One might try to respond to Bennett by saying that some effects are "too close" for one to be within the scope of the agent's intention and other outside of it. (Thus, if e1 and e2 are sufficiently 'close', then if X intends to bring about e1, then X intends to bring about e2.) But then what is needed is an account of when two effects are 'too close'. What we want is an account according to which civilian deaths and civilians blown to bits are too close but civilian deaths and the destruction of the armory (in STRATEGIC BOMBER) are not

¹⁰Foot, (1978)

too close. I turn now to a couple of attempts to give such an account. Again, these are attempts to say why it is unproblematic to claim the strategic bomber intends to produce the armory destruction but not the deaths while it is problematic to say the terror bomber intends to blow civilians to bits but not to kill them.

Appeal to uncertainty

One might think that Bennett's objection fails to take account of a difference in probabilities. Since intentions are dependent upon beliefs, the relevant notion would be some notion of subjective probability. According to this reply, the terror bomber is certain that if he blows the civilians to bits, they will die. On the other hand, the strategic bomber may believe there is a chance that he will cause destruction to the armory without causing civilian deaths.

I think this objection misses the point. As Bennett points out, "we have the notion of what is not intended but is foreseen as an *inevitable* by-product of one's means."¹¹ Defenders of the Doctrine of Double Effect do not have in mind that the difference in the bomber's intentions is traceable to any difference in probabilities. We can assume that the strategic bomber does not think there is a chance that the armory will be destroyed but that there will be no civilians deaths caused by his bomb--he considers the civilian deaths an *inevitable* upshot of dropping his bomb. Still, we want to say he does not intend to bring them about. (Perhaps both bombers believe there is a chance the bomb will not detonate, but this offers no grounds for treating them differently.)

¹¹Bennett (1981), 112

Appeal to Invariability

A second attempt to say when two effects are "too close" focuses on the following: even though the strategic bomber is certain that bombing this particular area at this particular time will cause civilian deaths, he believes that many armory bombings don't cause civilian deaths. It is often possible to blow up a military target without harming civilians. Not so for terror bombings. Blowing people to bits invariably causes death.

Warren Quinn has argued that appeals to invariability fail. He points out that hysterectomies do not invariably lead to the death of a fetus; in fact, presumably they usually do not involve any such death. So far, by the invariability criterion, a doctor might foresee that performing a hysterectomy in a certain case (the woman is pregnant) will cause a death but the doctor need not intend to bring about the death. However, notes Quinn, all hysterectomies performed on pregnant women do involve the death of a fetus. We can make trouble for the invariability criterion in the following way. Quinn asks us to suppose "that hysterectomies performed on patients who are in the early months of pregnancy are distinguished by the use of a special anesthetic that is safer for the patient and, in itself, harmless to the fetus." If this were so, then these hysterectomies could be separated out from other hysterectomies and one feature of these hysterectomies is that they would invariably result in the death of the fetus. Now it would turn out by the invariability criterion, that when a hysterectomy on a pregnant woman is performed, the foreseen death of the fetus is intended. Presumably the same sort of move could be made in the case of strategic bombing too. So the invariability criterion will not be helpful.

These last two replies attempt to say when two effects are 'too close' for one to be within the scope of the agent's intention, the other not. The following replies attempt in other ways to block Bennett's conclusion.

Objection to F-for-a-while

Consider the Disjunction Principle again. First there is the question of when feature F is a conjunction of features G and H. 'Being a bachelor' may well be a conjunction of 'being single' and 'being a male'. But we might wonder: Is 'being a yellow yardstick' a conjunction of 'being a stick which is yellow for eighteen inches' and then 'being a stick which is yellow for a second eighteen inches'? Is getting four hours of relief from cold symptoms a conjunction of getting one hour of cold relief and then three more? Let's assume so. Then it will follow that as I approach the yardstick with my paintbrush, I intend to bring it about that the yardstick be yellow only if I intend to bring it about that the first eighteen inches is yellow and I intend to bring it about that the second eighteen inches is yellow. Likewise, in taking some medicine, I intend to get four hours of relief from cold symptoms only if I intend to get the first hour of cold relief and I intend to get the next three hours of cold relief (and of course: only if I intend to get the first ten minutes of cold relief and then the next 3 hours and 50 minutes of cold relief, and so on). This all begins to look very unappealing. Furthermore, Bennett's troublesome results all involve cases of permanent and irreversible change. They all involve the claim that a feature of the form 'F forever' is a conjunction of the features 'F for a while' and 'F thereafter'. I will say that such cases involve 'temporal disjunction'. Why not simply disallow temporal disjunction?

Bennett's reply is that if we rule out such cases we disallow the reasonable conclusion that the politician intends to bring about only a one-month break up of the union. The politician is assumed not to intend to bring about the January stage of the union's inoperativeness since he is not motivated by the belief that he will bring that stage of inoperativeness about. But Bennett relies upon temporal disjunction to get from the claim that he January stage of inoperativeness is unintended to the claim that the permanent inoperativeness is unintended.

Ahistorical Reply

Someone might expect his action to produce an F and a G and have only the coarse-grained thought of an F-and-G, and not thought of producing an 'F without a G'. Gilbert Harman has raised this point and Bennett concedes that in such a case the test would seem to too finely delimit the scope of the person's intention. In certain cases, it will attribute to a person the intention to produce an F without a G when the person has no such concept. But Bennett is quick to point out that this does not rebut his argument since the terror bomber does have the concept of bodily inoperativeness.

Appeal to Structural Asymmetry

One might think that Bennett's point misses a structural asymmetry between the case of the strategic bomber and that of the terror bomber. The terror bomber believes that his ends (terror and a drop in enemy morale) flow causally from the deaths themselves. The strategic bomber believes that his end (weakening the enemy militarily) flows causally from the destruction of the armory but not from the deaths. However, this response is question-

begging. Bennett would reply that the terror bomber does not believe the drop in morale flows from the deaths but from something narrower, from the inoperativeness of the civilians. (And he can stress again that death itself isn't required since the terror bomber's aim is compatible with the civilians being miraculously restored to life at some later time.)

Anscombe's Approach

Following Elizabeth Anscombe, we might adopt as a test for whether someone intends to bring about a certain effect asking them why they are bringing it about. If their response is of the form "to..." then they intend to bring it about. But if their response is of the form "it couldn't be helped", then they don't intend to bring it about.

Warren Quinn explains why this sort of approach fails. Bennett's objection just needs to be rephrased. The point is, the terror bomber may respond to the question "why are you bringing about deaths?" by saying "'the actual deaths can't be helped if I am to create the realistic appearance of death and destruction.'"¹²

It is unsurprising that appealing to Anscombe's test is unhelpful. As I have pointed out, Bennett relies on a reductive account of intentions, like Anscombe's. This raises the question of whether Bennett's objection is dependent upon one particular view of intentions. (3') after all is a claim about what the terror bomber intends and (3') is supported by appeal to the counterfactual test (which is an expression of an Anscombe-type view of intentions).

¹²W. Quinn, "Actions, Intentions, and Consequences: The Doctrine of Double Effect", (1989), 340

Bratman's Approach

Remember that Bennett has assumed an Anscombe-like view of intentions. We might wonder whether his argument isn't really a reductio of Anscombe's view. Perhaps the problem does not arise on Bratman's view. Bratman's view certainly offers more resources than does Anscombe's for addressing objections like Bennett's. This is largely due to the fact that his account incorporates reasoning-centered aspects of intentions, not merely volitional ones. But I think in the end Bennett's objection makes trouble for Bratman's view too. This is most apparent with endeavoring and screening-out incompatible options because these are made out counterfactually. Let's look at endeavoring again.

Bratman claims that if the terror bomber were to believe he had missed the civilians, he would consider his raid unsuccessful and bomb again. But Bennett can argue that this does not show the bomber intends to *kill* the civilians. If the civilians were to miraculously recover after the war, the terror bomber would not circle back and bomb again. Or, if the bomber were to find out later that the bomb he had dropped was a new experimental sort of bomb which causes inoperativeness for a while, but not death, he would consider his raid successful as long as the inoperativeness lasted long enough to cause terror. Such counterfactuals either need to be ruled out as somehow illegitimate, or we need to conclude that the terror bomber endeavors to make civilians inoperative for a while, but not forever.

This is the endeavoring component. But as we have seen, Bratman's view has further resources. By looking at what the agent would screen off as incompatible, I think we'll get much the same result as we do with endeavoring (this is because both are roles get explicated by appeal to counterfactuals). But let's look at means-end reasoning. Surely a standard

terror bomber simply reasons about how to kill the civilians, not to make them inoperative. Perhaps this helps, but it's unclear. Bratman does not address situations in which the roles laid out above pull in different directions. If endeavoring suggests an agent intends to bring about an N as a means and means-end reasoning suggests the agent intends to bring about an M, what does the agent intend? One might try to address Bennett's objection by stipulating that in cases where endeavoring and means-end reasoning pull apart, more weight should be given to means-end reasoning. But we want to know why this should be so.

Furthermore, we might wonder why actual means-end reasoning is relevant to intentions as opposed to how an agent would reason? If we take into consideration how an agent would act, not just how an agent does act, then why not look at how an agent would reason? For example, if we suggested to the terror bomber that inoperativeness for a while would suffice to bring about his end, then he might well be disposed to think about how to bring about inoperativeness for a while. (Again, notice this is not so with an ethnic cleansing bomber.) And, even if we restrict ourselves to the agent's actual reasoning, trouble can be made for Bratman's account.

Imagine a group of terror bombers have been to West Point where they learned to think not about causing deaths, but merely about causing terror by causing bodily inoperativeness for a while. Their means-end reasoning is about how to bring on inoperativeness which will bring about terror.¹³ The most expedient way is to use a bomb which they foresee will cause deaths, but this, they are quick to point out, is unintended. Unless we are prepared to say that these bombers really do intend to bring about mere inoperativeness, then

¹³This particular formulation of the problem was suggested to me by Bratman.

there is no point in developing the option of putting more weight on means-end reasoning. It seems to me that in such a case, our response is that these bombers are somehow deceiving themselves, not that they really do intend to bring about mere bodily inoperativeness. It looks as though Bennett's objection makes trouble for Bratman's account of intending as well.

Even in spite of the failure of these replies, I am not prepared to follow Bennett in concluding that the notion intended as a means cannot be given a firm theoretical grounding.¹⁴ However, we have seen that the problem Bennett raises is a serious one. Warren Quinn accepts Bennett's argument against the Ordinary View but tries to salvage something of the Doctrine of Double Effect by offering an alternative principle. This alternative is meant to do the normative work of the DDE while avoiding Bennett's objection. I will be returning to Bennett, but for now, let's suppose as Quinn does that Bennett's objection holds and have a look at Quinn's view.

¹⁴Tanner lectures, p. 113.

Chapter Three: Direct Agency: Quinn's Response to Bennett

1. Introduction

The objection raised by Bennett is a serious one--it is not as simply solved as first glances may suggest. Warren Quinn responds to Bennett's objection by offering a new distinction, one which Quinn claims does the moral work typically attributed to DDE but which is not subject to Bennett's objection.¹ Quinn suggests that we distinguish what he calls 'harmful direct agency' from what he calls 'harmful indirect agency'. And, he suggests a two-part principle as a sort of replacement for DDE (I will call his two-part principle DDE*). The second part of DDE* is parasitic on the first part, but the first part is meant to stand alone--it is not dependent upon the second part.

2. DDE* part-one

The first part of DDE* is as follows:

DDE* part-one: Harmful direct agency is morally worse, and thus harder to justify, than harmful indirect agency.²

¹ Warren Quinn, "Actions, Intentions and Consequences: The Doctrine of Double Effect", 1987.

²Notice that Quinn gives a comparative formulation. It says that certain types of agency are worse than other types, it does not say that this difference matters to permissibility. In chapter 1, I gave a noncomparative formulation of DDE. For the most part, the differences between the comparative and noncomparative formulations won't be relevant here. I do discuss the two formulations briefly toward the end of the chapter.

Since the principle relies on the notions of harmful direct agency and harmful indirect agency, I will discuss these in turn. Quinn defines *harmful direct agency* as "agency in which harm comes to some victims, at least in part, from the agent's deliberately involving [the victims] in something in order to further his purpose precisely by way of their being so involved...". This definition of harmful direct agency has three parts; that is there seem to be three conditions which taken together provide necessary and sufficient conditions for harmful direct agency. *X is a direct agent of harm to Y* if and only if:

- (i) X harms Y³
- (ii) X deliberately involves Y in a plan in order to further her (X's) own purpose
- (iii) The harm to Y is at least partly a result of X's deliberately involving Y in the plan.

Let's take a brief look at each of these conditions and see what work each does.

Condition (i) plays two roles: first, it rules out as harmful direct agency any case in which no one is harmed. Second, it rules out cases in which someone is harmed but not by anyone's agency. For example, if Jed is harmed by a hurricane, this won't count as harmful direct agency.⁴

Condition (ii) works to rule out cases where one person causes harm to another, but in a completely accidental way. It will, for example, rule out a case where a driver accidentally hits a pedestrian. Such a case counts as one of

³By this I mean that X is the agent of harm to Y. If Z pushes X into Y, then Z harms Y, but X does not harm Y.

⁴It might be thought that (i) is too strong; that Quinn is attempting to introduce a subtler notion of agency which requires only that (1) Y is in fact harmed and (2) that the harm to Y is connected in the right way to an intention of X's. I discuss this view shortly.

harmful agency, but the case does not count as harmful **direct** agency since the driver does not **deliberately** involve the pedestrian in a plan.⁵

Condition (iii) links together (i) and (ii). (i) and (ii) leave open the possibility that the agent causes harm to the victim and the agent is deliberately involving the victim in a plan, but that the two are unconnected. It rules out the following kind of case:

Adam is need of cash. He comes up with a plan to rob a local store. His plan involves setting up his friend Mike to take the blame should anything go wrong. Around the same time, he accidentally hits a pedestrian with his car. The pedestrian turns out to be Mike.

⁵(ii) also rules out cases in which one person involves another in a plan but not for advancement of her own purpose. However, I think we should not take 'for her own purpose' to mean 'in her own interest'. Quinn categorizes the agency of the terror bomber as harmful direct agency even though the terror bomber may be acting in the interest of his country, not himself. We can say that the terror bomber involves the civilians for his own purpose if we understand this to mean simply that he thinks involving them advances the interests of one of his ends. (So I am assuming Quinn allows that our ends are sometimes to do with what's in the interest of others). I think that (ii) is best understood as ruling out cases in which the agent has the victim's own interests in mind. For example, if I involve you in a scheme in order to keep you out of the house while a surprise party is planned for you, then, even if harm should come to you as a result of this plan, the case will not count as harmful direct agency since I am attempting to do something for you rather than to further my own purpose.

Someone might object that I make the surprise party my own purpose; I have made your interests my own. So even in this case, I do deliberately involve you in a plan to further my own ends. Because Quinn is interested in the disrespect which he thinks attaches to involving others in one's plans, I think such cases don't count as satisfying (ii). When one person makes another's interests his own, this is typically a sign of respect, not disrespect. I will put aside complicated cases where one person misjudges what is in the interest of others as well as questions about paternalism.

Quinn specifies that the agent deliberately involves the victim in order to further his own purpose *precisely by way of their [the victim] being so involved*. I have dropped this last part from (i)-(iii) because I cannot see what it adds. I cannot think of a case in which X deliberately involves Y in something in order to further his (X's) own purpose but does not plan to further his own purpose precisely by way of Y being involved. To involve Y in order to further a purpose is to plan to further the purpose precisely by way of involving Y.

It also rules out a case in where one person is both agent and victim; Quinn is not explicit on this but I think it is natural to understand his account as ruling out such cases.

This case meets conditions (i) and (ii). Adam harms Mike and Adam is deliberately involving Mike in a plan to further his own purpose. But the two are not connected. We are to assume in the above case that Adam's plan to involve Mike in the robbery is not related to his hitting Mike. (iii) rules out such cases.

So far, I have been discussing harmful direct agency. DDE* part-one relies on a distinction between harmful direct and harmful indirect agency. Now let's look at harmful indirect agency. Quinn defines "harmful indirect agency" as "agency in which either nothing is in that way intended for the victims or what is so intended does not contribute to their harm."⁶ What does he mean by "in that way intended..."? He is referring to the way in which harm is intended in cases of harmful direct agency. He is saying that in cases of harmful indirect agency either condition (iii) is not met (what is so intended does not contribute to the agent's harm) or both (ii) and (iii) are not met (nothing is in that way intended for the victims). (If (ii) is not met, (iii) will not be either.) Notice that condition (i) is met in cases of harmful indirect agency, otherwise the case would not be one of harmful agency. So in these cases, although the agent harms the victim, the agent fails to deliberately involve the victim in a plan or the agent does involve the victim in a plan but the harm is unrelated to this involvement.

The account I have given of Quinn's distinction strays some from the most natural reading of his words, so I want to briefly defend my account. The way Quinn explains harmful direct agency may give the impression he is introducing a subtler notion of agency, one which does not require condition

⁶Quinn, page 343.

(i).⁷ A first reading of Quinn's words suggests that if X intentionally involves Y in something and as a result of this involvement Y is harmed, then X is a direct agent of harm to Y. X need do nothing more than intentionally involve Y in something which turns out to be harmful. I think that, in the end, this is not the best way to interpret Quinn and I will briefly mention a couple of reasons to reject this interpretation. First, notice that on such an account X's intentional involvement of Y will play a double role: it will be what makes the case direct agency not indirect, *and* what makes the case one of agency at all. The problem with allowing one condition to play this double role emerges when one considers harmful indirect agency. Recall that indirect agency is defined as agency in which "either nothing is in that way intended for the victims or what is so intended does not contribute to their harm." Take the first disjunct. If nothing is 'in that way intended for the victims' and it is just such an intention which makes the case one of agency at all, then in the absence of the relevant intention we have not a case of harmful indirect agency but a case involving no agency at all. On this reading, DDE* distinguishes harmful direct agency from accidents and other events. This is surely not right for a principle meant to do the work of DDE.

Notice though that this point argues only against taking (ii) and (iii) (p. 59) as the sole necessary conditions for harmful direct agency. It does not argue against something weaker than (i) which might provide an additional necessary condition for harmful direct agency as well as a sufficient condition for harmful indirect agency.

The second reason to reject the the more natural reading of Quinn's distinction is that this reading renders DDE* subject to devastating

⁷This is in fact how Fischer, Ravizza, and Copp interpret Quinn. in Fischer et al (1993). I discuss their objections to Quinn in the final section of this chapter.

counterexamples. I discuss some of these in the final section of this chapter when I evaluate the objections to DDE* made by Fischer et al.

I believe the idea behind Quinn's *harmful direct agency* is not to introduce a subtler notion of agency, such that having a harmful intention which causally contributes to actual harm is enough for to make one the direct agent of harm. Rather, I think Quinn's insight is that it does not matter whether, for example, the terror bomber's aim is to kill, dismember, or merely blow apart the civilians. We should be focusing instead simply on the fact that the terror bomber intends to do something harmful to the civilians as a way to further his end. Quinn is attempting to offer a principle which will mark the terror bomber's agency as morally impermissible whether or not the terror bomber is intending to cause death or mere inoperativeness.

Even if we accept (i) through (iii) as providing necessary conditions for harmful direct agency, DDE* differs importantly from DDE. Recall that DDE says:

The Doctrine of Double Effect (DDE): Where an agent, X, expects to produce two effects by A-ing, one of them bad and the other good, it is impermissible for X to A if she intends to bring about the bad effect and permissible for her to A if she does not intend to bring about the bad effect.

DDE divides cases of foreseen harm into two classes--those in which the agent intends to produce the harm and those in which the agent does not intend to produce the harm. DDE*, on the other hand, does not restrict itself to harms which are foreseen and most surprisingly, discriminates against

some cases in which the agent not only does not intend the harm but does not even foresee it. To see the difference, consider the following case:

Frank is having trouble with his car. He invites his friend Jane over, knowing he can persuade her to help him fix his car. While Jane is under the car, she asks Frank to turn on the engine and put the car in neutral. Frank accidentally puts the car in reverse and unwittingly backs it over Jane.

This case counts as one of harmful direct agency. Frank's agency satisfies conditions (i), (ii) and (iii). Frank harms Jane (by running the car over her). He deliberately involved her in a plan for his own purposes. And, she is harmed as a result of being involved in this plan. DDE* discriminates against this case while it falls outside the scope of the Doctrine of Double Effect.

I believe Quinn intended his principle to apply only to cases of foreseen harms; there are at least three reasons for thinking this is so. First, charity demands it: if his principle is not so understood, it is subject to crippling counterexamples.⁸ Second, he is offering a replacement principle for DDE. If we do not understand his principle to apply only to cases of foreseen harm then it will turn out to differ from DDE too much to be considered a replacement for DDE. Third, in a footnote Quinn suggests the doctrine of double effect (however it is formulated) puts a set of necessary conditions on morally permissible agency in which a foreseen harmful upshot is brought about. He suggests that his version of the doctrine works in the same way.⁹ I see no reason why Quinn could not have counted as harmful direct agency only those cases in which the agent foresees the harm. While it is

⁸I discuss such examples in the final section of this chapter.

⁹See footnote 3 of his article.

unfortunate that he didn't make this qualification explicit, I think he simply assumed it and that we should amend the above account of harmful direct agency so that it takes this qualification into account. I suggest we amend the conditions above for harmful direct and harmful indirect agency by replacing:

(i) X harms Y

with

(i') X harms Y and foresees that she will harm Y.

I take the following to be necessary and sufficient conditions for harmful direct agency:

(i') X harms Y and foresees that she will harm Y.

(ii) X deliberately involves Y in a plan in order to further her (X's) own purpose

(iii) The harm to Y is at least partly a result of X's deliberately involving Y in the plan.

With this change, DDE* will now divide cases of foreseen harm into two classes—direct and indirect--and thus will be more like DDE in an important respect. Here is a reminder of DDE* part-one:

DDE* part-one: It is morally worse to bring about harm by direct agency than it is to bring about harm by indirect agency.

The second part of Quinn's replacement principle relies upon a distinction of harmful direct agency into two types: opportunistic agency and eliminative agency. I will discuss this further division in more detail later. For now, a statement of DDE* part-two is as follows:

DDE* part-two: opportunistic harmful direct agency is morally worse than eliminative harmful direct agency.

Quinn believes that there are three things favoring DDE* over DDE. First, he thinks that because DDE* employs a distinction which differs slightly from the distinction between intended and merely foreseen consequences, it is not plagued by Bennett's objection.¹⁰ Second, he thinks that DDE* can explain more about our intuitions in paired cases than DDE. Specifically what he has in mind is this: he thinks that intuition tells us the terror bomber's action is morally worse than the strategic bomber's. Likewise says Quinn, intuition tells us that a surgeon who performs a craniotomy which results in the death of the fetus does something morally worse than a surgeon who performs a hysterectomy on a pregnant woman, where this also results in the death of the fetus. But Quinn believes that we intuitively find the moral gap between the bombers larger than the gap between the surgeons. Furthermore, he believes that this asymmetry cannot be explained by DDE but can be explained by DDE* (in virtue of its two parts working together). Third, Quinn claims that it is easier to see the rationale for the distinction drawn by DDE* than it is to see the moral significance of the distinction relied upon by DDE.

In the rest of this chapter, I argue that although DDE* appears to avoid Bennett's objection, in fact it does not. Bennett's objection applies to DDE* as well. Then I look at Quinn's claim that DDE* can explain asymmetries in our intuitions about the different pairs of cases. Here I examine DDE* part-two and argue that it does not divide up some of the relevant pairs of cases in the

¹⁰Chapter 2 contains an extensive discussion of this problem

right way. Thus, two of the points DDE* is taken to have in its favor over DDE do not hold up. Finally I defend DDE* part-one from some recent criticism in the literature. I conclude that we should reject DDE* part-two and that DDE* part-one and DDE are roughly on a par--neither helps get around Bennett's objection.

3. DDE* Applied to the Bombers

Quinn claims that DDE* draws a moral distinction between the bombers in a way which is not subject to Bennett's objection, or to problems of "closeness" generally. To see what Quinn has in mind, consider the terror bomber first. Of the terror bomber, Quinn writes:

...the bomber undeniably intends in the strictest sense that the civilians be involved in a certain explosion, which he produces, precisely because their involvement in it serves his goal... [H]is purpose requires at least this--that they be violently impacted by the explosion of his bombs.¹¹

Quinn's point is that even if the terror bomber does not intend to cause the actual deaths of the civilians, he at least intends to involve them in something which is harmful. And this, it turns out, will suffice for harmful direct agency. The case meets the necessary and sufficient conditions for harmful direct agency set out above. The terror bomber harms the civilians. The terror bomber deliberately involves the civilians in his plan, and the harm comes to the civilians at least in part by way of the bomber deliberately involving them. Notice the harm itself need not be intended; what matters is that the involvement of the civilians is intended. This then gets around the

¹¹Quinn, p. 342.

question of whether the bomber intended to cause the civilians deaths, or merely their dismemberment, or perhaps only their temporary inoperativeness.

Now let's turn to the strategic bomber. Quinn argues that this is a case of harmful indirect agency. He reasons as follows: STRATEGIC BOMBER is a case of harmful agency since condition (i) is met--the strategic bomber harms the civilians and foresees he will do. But it is not a case of harmful direct agency since condition (ii) is not met. The strategic bomber does not deliberately involve the civilians in his plan. Their involvement is incidental to his purpose.

This then is what I think Quinn has in mind when he claims that DDE* is not subject to Bennett's objection. Bennett's objection was to the notion 'intended as a means' and DDE* does not rely on this notion.

But I think Quinn is wrong. The problem is with whether condition (ii) is met. Quinn claims that that the strategic bomber's agency is not direct because condition (ii) is not met: the strategic bomber does not deliberately involve the civilians in his plan. But why think this? After all, the strategic bomber knows the civilians are present, he deliberates over what to do. He chooses to bomb, foreseeing he will cause them harm. Whether or not this is a case of harmful indirect agency hinges on whether condition (ii) is met. Quinn uses 'deliberately' and 'intentionally' interchangeably. I can find no reason to think he has in mind an account of doing something deliberately which differs importantly from doing it intentionally or intending to do it. So the question remains: why think the strategic bomber does not deliberately involve the civilians; otherwise put, why think he does not intend to involve them?

One might try to defend Quinn's assumption that (ii) is not met by noting that the strategic bomber would answer No to the following counterfactual test question: "If you believed that no civilians would be involved in the explosion, would you be less likely to bomb?" After all, the strategic bomber would be just as happy were no civilians around. His plan does not in any way require the presence of civilians. Notice by way of contrast that the terror bomber would answer Yes. He would change his plans if he were to find no civilians were in the area.

So far, it looks as though Quinn is right that DDE* avoids Bennett's objection. But I think that DDE*'s success at dividing up the bombers the right way is dependent on the particular way in which the terror bomber case happens to be described. To show this, I will introduce an altered version of the terror bomber case which I think will cause problems for DDE*.

4. City Bomber Objection to DDE*

Consider the following case:

CITY BOMBER: The city bomber wishes to demoralize enemy leaders. He knows that the leaders are residing in a hideaway in the mountains above the city. Here they have a view of the city and are free from danger. The city bomber knows that if he explodes bombs over the center of the city, the leaders will be convinced that the city population has been destroyed and are likely to concede.¹²

¹²This case is slightly varied from a case Quinn discusses (see f.n. 16) and attributes to David Lewis. Quinn argues that DDE* handles Lewis's case but I think DDE* will not handle this slightly varied case.

Now I want to argue (1) that the city bomber's harmful agency is indirect, not direct, but (2) the city bomber is morally closer to the terror bomber than the strategic bomber. So I am arguing that DDE* should discriminate against the city bomber, just like it discriminates against the terror bomber, while in fact it discriminates in favor of the city bomber. Another way of putting this is that DDE* draws a line which leaves the terror bomber on one side and both the city bomber and the strategic bomber on the other side. I think it should draw a line between which puts the city bomber on the same side as the terror bomber.

The city bomber believes that to achieve his purpose, he must target an area which the leaders believe to be heavily populated; the city center serves this purpose. The city bomber's agency is direct only if he deliberately involves the population in the explosion. Does the city bomber deliberately involve the civilians in his plan? Let's appeal again to a counterfactual test question. We might ask the city bomber: "If you had believed that no civilians would be involved in the explosion, would you have altered your plans?" The city bomber will answer No. He will answer No because he is indifferent to whether there are any civilians actually in the city center at the time he bombs. He believes that to achieve his purpose he needs only to cause the leaders to believe that civilians have been killed. And this can be done by dropping a bomb or bombs over the city center regardless of whether civilians are present or not.

Just as the terror bomber was indifferent to the actual deaths of civilians (he needed only to cause a belief in the deaths), the city bomber is indifferent to the actual involvement of the civilians (he needs only to cause a belief that civilians were involved). If the city bomber came to believe that

civilians had been warned of the upcoming bombing and (it being a fairly small city) had all evacuated in the night, he would bomb just the same. As long as he thinks the city leaders believe the city center is populated, then whether it is actually populated or not is irrelevant to his purpose.

This example should make clear that DDE* does not escape Bennett's objection. Now, instead of a problem about whether the terror bomber intends to cause deaths, or merely temporary inoperativeness, we have a problem about whether the city bomber intends to involve the civilians themselves or only the city center. This new problem arises out of the same sorts of considerations as the old one (just as the terror bomber does not require actual civilian deaths, the city bomber does not require that any civilians be actually involved.)

Quinn seems to have thought that Bennett's objection could be avoided by looking not at whether an agent intends to cause harm but at whether the agent intends to involve someone in a plan. But Bennett's objection is to the notion 'intend as a means'. While this notion plays a more visible role in DDE (which rests upon the distinction between intended and merely foreseen harms), it plays a crucial role in DDE* nonetheless.¹³ Since Bennett's objection is to this notion, switching the focus of the principle from whether certain an agent intends to harm as a means (DDE) to whether the agent intends to involve the victim... as a means (DDE*) will not help.

My objection to DDE* relies upon the assumption that DDE* ought to discriminate against the city bomber. I assume that if the original terror bomber case counts as harmful direct agency, the city bomber case ought to

¹³As I have noted above (fn 4) Quinn is not trading on a difference between 'deliberately does A and 'intends to do A. This can be seen by comparing his definitions 'harmful direct agency' and 'harmful indirect agency'. He uses one expression in the former and the other in the latter.

count as harmful direct agency too. But there is room for disagreement here. Quinn might argue that it is a point in favor of DDE* that it treats the terror bomber and the city bomber differently. After all, the former targets people while the latter targets places. This, he might argue, makes the city bomber's agency morally closer to that of the strategic bomber than to that of the terror bomber.¹⁴

I think this reply is inadequate. If the city bomber himself issued a warning to the civilians so that they could, unbeknownst to the leaders, evacuate in the night, then I think this response would be right; there is an important moral difference between that city bomber and the city bomber I discuss. The city bomber I discuss does not do this. He chooses the city center because it is an area which he thinks the leaders will assume to be populated (at the time of the bombing). Let us assume that he plans to drop his bomb on Wednesday evening. Then, he will look for an area which the leaders have every reason to believe is populated at Wednesday evenings. Again, if he were to search and search for an area such that it is usually populated Wednesday nights and which is such that the leaders will think it is populated on the Wednesday night of his bombing, but which in fact, for some reason or other, is not populated on this particular Wednesday evening, then morally this bomber would be more like the strategic bomber. But he does no such thing.

The force of Bennett's objection is this: although the area the city bomber chooses is populated, the city bomber believes that the area need not

¹⁴In fact, I think Quinn would not have argued this way. See fn 16 of Quinn's paper. Quinn explains that he altered his original formulation of harmful direct agency in response to an objection posed by David Lewis. It seems to me that if Quinn were going to argue here that targeting city centers is importantly different from targeting the population, he would have made a similar response to Lewis instead of changing the formulation of 'harmful direct agency'.

be currently populated in order for him to achieve his end. Now I want to claim that this fact about him does not make him morally closer to the strategic bomber than the terror bomber.

This point may come out more clearly if we consider an actual case of city bombing-the United States bombing of Hiroshima in World War II. The aim of the U.S. was to prove to the Japanese the extent of American power. Let's charitably assume the following: Had it been possible to drop an atomic bomb over Hiroshima, cause the destruction of the city, and yet not cause harm to civilians, Truman and the U.S. military leaders would have chosen to do so. Even if this is so, surely it does not make the actual bombing a strategic bombing rather than a terror bombing.

For someone to argue that the city bomber case is morally on a par with the case of the strategic bomber, not the terror bomber, he would have to be willing to morally distinguish bombers who deliberately involve people in their plans, (that is, they require that the area they target is currently populated) from bombers who do not deliberately involve people in their plans but do deliberately involve areas which they know to be populated, like hospitals, city centers, and malls. Furthermore, certain bombings (like city bomber's) would no longer count as terror bombings, even though they work (reach their end) by creating terror.¹⁵ I think this is an undesirable route to take and I think it is a route Quinn would not want to take.

With regard to Bennett's original objection (that the terror bomber does not intend to cause deaths, just temporary inoperativeness), Quinn writes: "This line of objection clearly threatens to deprive the doctrine [of double effect] of most of its natural applications." I think many proponents of

¹⁵This was pointed out to me by Walter Sinnott-Armstrong.

Double Effect would say the same here; if DDE* avoids Bennett's objection only by distinguishing terror bombers who bomb city centers and malls because they are currently populated from terror bombers who target these locations because the leaders believe them to be populated, then the doctrine has been deprived of most of its natural applications.

I have been assuming that a supporter of Quinn would appeal to a counterfactual test question to defend the claim that the strategic bomber does not intend to involve the civilians while the terror bomber does. But perhaps Quinn was rejecting Bennett's counterfactual test. Here I think supporters of Quinn face a dilemma. Either they provide us with an account of *intending to involve as a means* which does not rely on counterfactual test questions or they appeal to counterfactuals. If they appeal to counterfactuals, then, as I have just argued, they will run into Bennett's objection. If Quinn has some other account of the notion 'intend to involve as a means', then why switch introduce the notions 'direct agency' and 'indirect agency' at all? Why not stick with the original Doctrine of Double Effect?

My own view is that Quinn did not have in mind an alternative to the counterfactual test. Quinn's response to Bennett does not seem to be that we ought to reject counterfactual test questions, but that we ought to see that we didn't have quite the right distinction before. The distinction between intended and merely foreseen consequences is not quite the significant distinction, but the one between harmful direct agency and harmful indirect agency is the distinction which is morally significant. Quinn suggests that what is importantly different about his view is not that he has a different idea of intending, rather, it is that DDE* focuses on whether or not an agent intends to involve the victim, rather than on whether an agent intends to

cause the harm itself. But then, we do not have a way around Bennett's objection.

To summarize, Quinn's notion *direct agency* relies upon the notion *intending as a means*. Bennett has made trouble for the idea of intending as a means. Quinn must either (1) offer an unproblematic way of making out what an agent intends as a means, or (2) show that what was a problem for intending versus foreseeing is not a problem for direct versus indirect agency. If he has an unproblematic way of making out what is intended versus what is merely foreseen, then why introduce the new distinction at all? If he does not have an unproblematic way of making out what is intended, then he really can't get around Bennett's objection by introducing *direct agency*, since this notion too relies upon the notion *intended as a means*. In short, DDE* part-one does not avoid Bennett's objection. Now let's turn to DDE* part-two.

5. DDE* part-two: Opportunistic and Eliminative Agency

Recall that DDE* part-two says:

DDE* part-two: opportunistic harmful direct agency is morally worse than eliminative harmful direct agency.

Quinn's reasons for thinking a second part to DDE* is needed will be clearer if we start by considering another pair of cases.

HYSTERECTOMY: A pregnant woman will die unless her cancerous uterus is removed. If the operation is performed, given the limits of medical technology, the fetus will die. If the

operation is not performed, the woman will give birth to a healthy infant but will not survive the birth.

CRANIOTOMY: A pregnant woman will die unless the fetus she is delivering is removed from her body immediately and removing it immediately requires crushing its skull. There is another way to remove the fetus which will allow the fetus to live, but involves allowing the woman to die.

Quinn notes that there seems to be a much more pronounced moral gap between the terror bomber and strategic bomber cases than there is between CRANIOTOMY and HYSTERECTOMY. By this he means that he thinks we find CRANIOTOMY in some way morally worse than HYSTERECTOMY and we find the terror bomber case in some way morally worse than the strategic bomber case, but the moral gap between the bombers is bigger than the moral gap between the surgeons.¹⁶

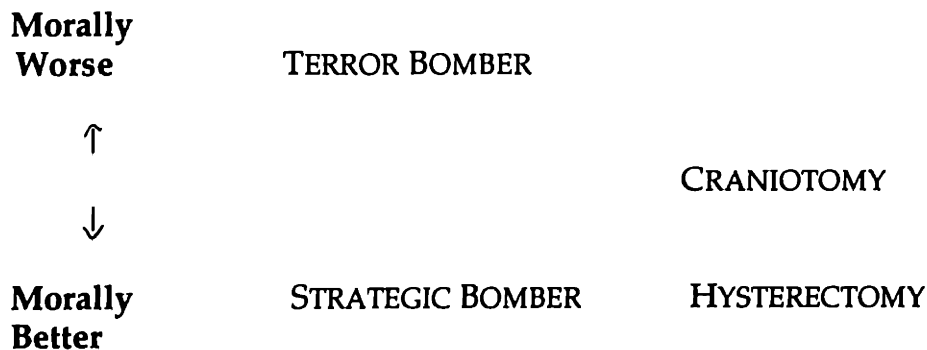
I should point out that intuitions regarding the surgery pair of cases are much less uniform than intuitions on the bomber cases. Also, as Quinn acknowledges, there are many features complicating the second pair of cases, for example, that the fetus is threatening the woman's life and that fetus is not yet a person and thus whether it has the same moral standing as the woman is extremely unclear. So it is no surprise that even those who do see a moral difference between CRANIOTOMY and HYSTERECTOMY are apt to find the difference less striking here than in the case of the bombers. Though

¹⁶ By speaking of one case as morally worse than another, I am avoiding a question about whether it is the action of one agent that is morally worse than the action of the other agent or whether the actions are morally on a par but one agent is more blameworthy than the other. I avoid this issue on purpose because I am appealing to intuition and I think intuition is not this fine-grained; rather our intuitions about the cases are responses to the whole situation.

Quinn acknowledges that such features may be responsible for the asymmetry in people's intuitions about the bombers versus the surgeons, he also maintains that the asymmetry may be explained by DDE* part-two.

Let's return to the asymmetry in intuitions. Perhaps the following chart will make Quinn's point clear.

Chart of Harmful Agency



On the chart, the gap between the two bomber cases is larger than the gap between the surgeon cases. This is supposed to represent a larger moral difference between TERROR BOMBER and STRATEGIC BOMBER is greater than between CRANIOTOMY and HYSTERECTOMY. Quinn suggests that this asymmetry shows a need for more line-drawing. He writes:

We have not yet defined the difference between the more pronounced moral asymmetry of... SB and TB [STRATEGIC BOMBER and TERROR BOMBER], and the apparently weaker asymmetry of HC and CC [HYSTERECTOMYand CRANIOTOMY]. This difference may partly depend on whether the agent, in his strategy sees the victim as an advantage or as a difficulty. In CC the doctor wants the fetus removed... Its presence there is the problem. Perhaps it would not be surprising if we regarded fatal or harmful exploitation as more difficult to justify than fatal or harmful elimination. If so, we might say that the doctrine

discriminates against direct agency that benefits from the presence of the victim (direct *opportunistic* agency) and more weakly discriminates against direct agency that aims to remove an obstacle or difficulty that the victim presents (direct *eliminative* agency)."¹⁷

So harmful direct agency is *opportunistic* if and only if the agent views the victim (or the victim's presence) as a benefit. And harmful direct agency is *eliminative* if and only if the agent views the victim (or the victim's presence) as an obstacle.¹⁸ But questions remain. For example: Is the distinction between opportunistic and eliminative agency exhaustive? Quinn says the asymmetrical judgments "may partly depend on whether the agent... sees the victim as an advantage or as a difficulty." This suggests an exhaustive division, but it's fairly weak evidence. Does an agent who involves another in a plan always view that other as either an advantage or a difficulty? Also, could one person view another as presenting both an obstacle and a benefit at once (with regard to the same plan)? I'll assume the two categories are mutually exclusive and the division exhaustive.

Also, we might wonder: how are we to understand eliminative direct agency? After all, Quinn tells us that in cases of direct agency, the agent deliberately involves the victim in something in order to further his purpose precisely by way of their being so involved." It sounds as though a feature of any case of direct agency is the expectation of the agent to benefit from the presence of the victim, since the agent involves the victim in order to further his plan. For agency to be both direct and eliminative, the agent must both deliberately involve the victim in a plan in order to further some purpose

¹⁷Quinn, p. 344

¹⁸I have made out these notions by appeal to the agent's view of the victim. This may not be evident from the quotation since Quinn says: "agency which benefits from..." But elsewhere Quinn focuses on the difference in how the agent views the victim.

and at the same time view the victim as presenting an obstacle. Perhaps we can best see what Quinn has in mind by considering CRANIOTOMY since he tells us this is a case of eliminative agency.

In CRANIOTOMY the surgeon does view the fetus as an obstacle and she also does deliberately involve the fetus in her plan. But she does not benefit from the presence of the fetus since her whole plan just is to eliminate it. Perhaps the following works as a test for eliminative agency:

Harmful Direct Agency is eliminative if and only if the agent's end is achieved by the victim's going out of existence.

If the fetus were to suddenly go out of existence, the surgeon's end would be achieved. However, consider this variant of CRANIOTOMY.

CRANIOTOMY TWO: A pregnant woman with a cancerous uterus will die unless both the fetus and the tumor are removed from her body immediately. The only way to save the woman is for the surgeon to first remove the fetus, by crushing its skull, and then remove the tumor.

CRANIOTOMY TWO is surely still a case of eliminative agency, rather than opportunistic. The only difference between the cases is the addition of the tumor; the surgeon's attitude toward the patient hasn't changed. But the above account of eliminative agency misses this case; it misses it simply because the surgeon must do more than remove the fetus in order to accomplish her end. Perhaps then the following revised account of eliminative agency is more plausible:

Harmful Direct Agency is eliminative if and only if the agent's goal is either advanced or achieved by the victim's going out of existence.

This new account is sufficiently weakened to capture CRANIOTOMY TWO. Let's assume that this account captures all the cases of eliminative agency. Does it exclude all cases of opportunistic agency? It will exclude most of them, and I think that the account captures the spirit of the distinction. Unfortunately however, there is a certain class of cases it seems to misclassify. TERROR BOMBER falls into this category. The terror bomber views the civilians as providing a benefit, not an obstacle. This is surely supposed to be a case of opportunistic agency. But the civilians provide a benefit precisely because their presence allows him to kill them, and thus to cause terror. Notice that if the civilians were to go out of existence, the terror bomber's goal might well be advanced. One might try to patch up the criterion in the following way:

Harmful Direct Agency is eliminative if and only if the agent's goal is advanced or achieved by the victim's going out of existence before the agent acts.

This last qualification is meant to help in cases like TERROR BOMBER. If civilians were to go out of existence before the terror bomber even began his raid, then their going out of existence would not advance his plan. But even this formulation would need to be modified. To see why, assume that the relevant action of the terror bomber's is dropping a bomb. If, just as the terror

bomber flies overhead but before he releases a bomb, civilians go out of existence, the enemy may believe that the bomber is responsible for the disappearance of the civilians (perhaps via some new very powerful sort of bomb) and so the enemy may still be terrorized.

I think rather than adjust the criterion above to accommodate such cases, we should consider the following question: Both the terror bomber's and the surgeon's goals are advanced by their victims going out of existence. What then makes one's agency opportunistic and the other eliminative? In answering this, I think we ultimately circle back to where we started: the surgeon sees the fetus as an obstacle while the terror bomber sees the civilians as providing a benefit. So in the remainder of the discussion I will stick with Quinn's way of making out the distinction, though it is somewhat imprecise.

We have been considering the difference between opportunistic and eliminative agency. But all of this is by way of examining the second part of DDE*. We might put DDE* part-two as follows:

DDE* part-two: opportunistic harmful direct agency is morally worse than eliminative harmful direct agency.

DDE* part-two is supposed to capture the asymmetry in the cases above as follows. In HYSTERECTOMY, the fetus is harmed by way of **indirect** agency whereas in CRANIOTOMY, it is harmed by way of **eliminative direct** agency. So we have a case of eliminative direct agency contrasted with a case of indirect agency. With the bombers, we have a case of **opportunistic direct** harmful agency (TERROR BOMBER) contrasted with a case of harmful **indirect** agency (STRATEGIC BOMBER). Below is a revised version of the Chart of Harmful Agency which includes Quinn's categories of Harmful Agency.

Chart of Harmful Agency

Direct Opportunistic Agency	TERROR BOMBER	
Direct Eliminative Agency		CRANIOTOMY
<hr/>		
Indirect Agency	STRATEGIC BOMBER	HYSTERECTOMY

This difference, according to Quinn, may explain the asymmetry in our moral judgements of the cases. We should expect DDE* to discriminate strongly in favor of the strategic bomber. And since the second pair of cases contrasts an instance of eliminative harmful direct agency (CRANIOTOMY) with one of harmful indirect agency (HYSTERECTOMY), we should expect DDE* to discriminate in favor of HYSTERECTOMY, but only weakly.

But now I want to turn to two paired cases which I think cause trouble for DDE* part-two. The first pair involves trolleys, the second, politicians.

6. Objections to DDE* part-two

To see the trolley cases, we need to start with Philippa Foot's well-known problem of the runaway trolley. The problem is this:

TROLLEY: A driver of a runaway trolley sees he is headed for five people who are trapped on a narrow track ahead. There is a side-track to the left, but there is a person trapped there as well. There is no way for any of these people to escape, so if the train

goes straight, five will be killed. If he steers the train to the left-hand track, one will be killed.

And the second case:

TRANSPLANT: A surgeon has just been brought five accident victims, each with serious damage to a different vital organ. All are in imminent danger; the surgeon's only hope of saving the five is to cut up a healthy patient who is in for a routine physical and distribute the organs to the five in need. We are to assume that the surgeon is highly skilled and success is certain.

In each case, the agent must choose between the lives of five and the life of one. Most people strongly believe that the trolley driver may turn the trolley to the left, causing the death of one rather than five, but that the surgeon may not cut up the healthy person to save her five. The problem is to say why the trolley driver may turn his trolley, killing one, but the surgeon may not operate, killing one.

So far, DDE* can offer an explanation. If the surgeon were to cut up the one against his will, her agency would be direct and opportunistic whereas if the trolley driver turns his trolley to the left, his agency is indirect. To see this, we need to appeal to conditions (i')-(iii) (p. 65). I'll start with TRANSPLANT. Condition (i') is met: the doctor harms the healthy person and foresees she will. (ii) is met as well. The surgeon deliberately involves the healthy patient in a plan. The relevant counterfactual test question is this: "If you had believed that the healthy patient would not be involved in the operation, would you have altered your plans?" It is an awkward question; it's awkward

to imagine the surgeon thinking over whether she would operate if the victim were not to be involved in the operation! While the question is awkward, I think it is clear that the surgeon's plan is completely dependent upon her belief that the patient will be involved. The surgeon would answer Yes.

Condition (iii) is met too. The harm which results to the victim is certainly a result of the surgeon's deliberate involvement of the victim in her plan. So this is a case of harmful direct agency. Furthermore, it is opportunistic agency. The surgeon views the victim as a benefit since without him she would be unable to achieve her end.

Now let's turn to TROLLEY. As I said above, I think this is a case of harmful indirect agency. If the driver turns the trolley, he harms the one on the track.¹⁹ But does the agent, the driver, deliberately involve the victim in his plan? I think he does not. Again, let's turn to a counterfactual test question: "If you had believed that the person on the left-hand track would not be involved in any way by your turning the train, would you have altered your plans?" The driver would answer No. He, like the strategic bomber, would be happy to hear that the victim would not be involved. Notice that his agency is like the strategic bomber's in that (1), unlike the terror bomber, his end is not jeopardized by the victim going out of existence and (2), unlike the craniotomy surgeon, his end is not achieved by the victim going out of existence. So the presence of the person on the left-hand track is incidental to his plan just as the presence of the civilians is incidental to the plan of the strategic bomber.

¹⁹One might claim that the driver does not cause the harm; the out of control trolley does. But surely the driver's turning the trolley left is a cause of the harm. Notice that if no one were on the track straight ahead and the driver still turned the trolley left onto one, we would not hesitate to attribute agency to him.

Now to the problem case. There is a variant on the trolley case which I think causes problems for Quinn. In the variant case, introduced by Judith Jarvis Thomson, the trolley tracks no longer simply diverge.

LOOP TROLLEY: As in TROLLEY, the driver of a runaway trolley faces a choice of allowing the trolley to go straight ahead where five are trapped on the tracks or to turn the trolley to the left where one is trapped. But here, the the left-hand and right-hand track curve around to form a loop, so if the trolley driver does nothing, the trolley will continue straight ahead, killing the five on the tracks and coming to a stop before rounding the bend to the one. If the driver turns the train to the left, the one will be killed, but because the one is large, his body will stop the train from continuing around the loop and hitting the five.

If the driver turns the trolley to the left, his agency will be opportunistic harmful direct agency. He foresees that the one on the left-hand track will be harmed. But does the driver deliberately involve the one in a plan? It seems that he does. Again, let's appeal to a counterfactual test question: "If you had believed that the person on the left-hand track would not be involved if you turned the trolley, would you have altered your plans?" If the driver believed that the one on the left-hand track would not be involved, or effected, by his turning the trolley, then he would alter his plans. He turns the trolley to the left because he sees a way to stop it from hitting five. If he were to come to believe that the person in the left would no way be involved, then he would no longer plan to turn his trolley left, so he would alter his plans. This shows the driver intentionally involves the one on the

left in his plan. Furthermore, conditions (iii) is met. The harm to the one on the left comes about as a result of the driver involving him in a plan. So Loop TROLLEY is a case of harmful direct agency.

LOOP TROLLEY causes problems for DDE, so if it causes problems for DDE*, this at least is not a strike in favor of DDE. Still, given Quinn's two-tier distinction, we might have thought that LOOP TROLLEY would at least fall into the eliminative rather than opportunistic category. It does not. The driver in LOOP TROLLEY views the one on the tracks much the way the surgeon in TRANSPLANT views her victim. The one provides a benefit, an opportunity for the driver to save the five. So LOOP TROLLEY counts as opportunistic agency. If the one the left hand track were to suddenly vanish, the driver's end would not be achieved, and in fact, he would be unable to achieve his end. DDE* puts these two cases far apart—it likens the gap between them to that between the two bombers, not to that between the two surgeons.

This case presents a problem for Quinn because many people think that the driver may turn his trolley left in both TROLLEY and in LOOP TROLLEY, yet according to DDE*, the driver's agency in LOOP TROLLEY is morally worse than the driver's agency in TROLLEY.²⁰ There may of course be those who distinguish two types of drivers in loop cases: the drivers who were going to turn left anyway (in spite of the presence of the one) and the drivers who turn left because of the one. I take it that Thomson meant for the driver in her LOOP TROLLEY example to be of the second kind and again, I take it that many people have the intuition that the driver may still turn his trolley

²⁰This pair of cases also speaks against DDE* (part-one & part-two). DDE* puts the moral gap between these two cases on a par with the gap between the bombers not the gap between the surgeons.

(and that the driver does not do something morally worse in this case than in the original trolley case).²¹

While I think that this pair of cases causes trouble for Quinn, intuitions on these cases are not uniform--some people believe that what the driver in LOOP TROLLEY does is morally worse.²² While this is not an intuition I share, I think this pair of cases is especially troublesome--it's hard to know what to say about it. So I will turn to a final pair of cases, which I think causes a more clearcut problem for DDE*.

The following cases require us to imagine that it is once again election time. Consider these politicians:

Max, a young, inexperienced candidate aims to get elected to Congress. He knows that Lois, a liberal senator from another district, is particularly disliked by the more conservative people of his district. He believes that he can get himself reelected by discrediting Lois. To discredit her, he claims that she is supporting a crazy economic plan. His scheme works and he is elected by his district while she is not by hers.

This is a case of harmful direct opportunistic agency. I think it is clear that Max harms Lois's interests.²³ Max deliberately involves Lois in his plan. If Max were to come to believe that his plan did not involve Lois, he would

²¹I mention this because one might hold the view that it is permissible for each driver to turn his trolley but still hold that the driver's agency in Loop Trolley is worse. Such moral intuitions might not be problematic for Quinn since DDE* says nothing about permissibility.

²²For example, Michael J. Costa, (1987).

²³The DDE is often applied to grave cases of harm -- like death or bodily injury. I assume this sort of harm--harm to one's interests--does count. I suspect that if this is not right, cases can be concocted which would make the same point and would count as cases of harm.

alter his plans. Furthermore, the harm comes to Lois by way of this involvement. Finally, Max sees Lois as presenting a benefit; he does not see her as an obstacle. Her existence or presence gives presents him with an opportunity to achieve his goal.

Now on to the second case:

Fred has political aspirations. He longs to represent his district in Congress. But Lois has held the seat Fred wants for years and years. Fred decides that it's high time to discredit Lois so that he himself can be elected. He devises a hair brained economics scheme and runs a series of ads in which he says that Lois supports this scheme. As a result of these ads, Lois is not re-elected.

Fred's agency, like Max's is of the harmful direct variety. Furthermore, Fred deliberately involves her in a plan of his. If he were to come to think his plan would not involve Lois, he would alter it. Finally, the harm which comes to Lois comes by way of her involvement in this plan. She is not re-elected because of Fred's scheme. But this case, unlike the previous one, is a case of eliminative direct agency, not opportunistic agency. Fred, after all, sees Lois as an obstacle not as presenting a benefit.

DDE* will discriminate against both of these cases, but it will discriminate more weakly against Fred's agency, since it is only eliminative harmful direct agency, not opportunistic agency. However, I cannot find any intuitive basis for thinking that Max's agency is morally worse than Fred's.

If we reject DDE* part-two, what are we to say about the asymmetry between the bombers on the one hand and the surgeons on the other? I think there is much we can say to explain the asymmetry. It may be traceable, for example, to a doubt about whether the fetus has rights (we have no such doubts about the civilians in STRATEGIC BOMBER), or perhaps we think that the doctor may feel a more sincere sort of remorse than the bomber. These cases, I believe, suggest that Quinn is wrong to think that the distinction between opportunistic and eliminative agency tracks our intuitions.

7. Review of Fischer, Ravizza, Copp

In a recent paper, John Fischer, Mark Ravizza, and David Copp appear to make much trouble for DDE*. However, I think they focus on many cases which are not in fact problematic for Quinn and do not focus on the real problem for Quinn. As I have already noted, at least some of their objections depend upon what I think is a mistaken, or at least uncharitable reading of Quinn. The real problem, I think, is that Quinn has motivated the introduction of his distinction by saying it gets around Bennett's objection--and it doesn't. If Bennett's objection stands, DDE and DDE* are both in trouble. Still, if DDE* can be defended against the objections of Fischer et al , then it is at least left roughly on a par with DDE. On the other hand, if Fischer et al are right, then DDE* does a much worse job than DDE of dividing up paired cases in a way which matches intuitions. I want to argue that Fischer et al are wrong. I will be suggesting that most of their objections rest on a misinterpretation of Quinn and that the remainder of their objections are simply mistaken.

The objections raised by Fischer et al fall into three types. I will give a sample of each type of objection and say how in each case I think DDE* can be defended. The first type of objection exploits Quinn's failure to make explicit that DDE* is restricted to cases involving foreseen harms. Of the objections presented, it is the ones in this group which initially appear most devastating. Here's one of their examples:

... Jack plans to ask Mary out for a date. He knows Mary will be in a certain area of the library at a particular time. Just before that time, Jack sees Sam there. He realizes he would feel awkward if Sam were sitting there while he asked Mary for the date.... In order to persuade Sam to go to the coffee house, he tells Sam a concocted story about some music being played there.... Sam takes the bait and drives toward the coffeehouse. Unfortunately, as Sam is proceeding toward the coffee-house, he is involved in an accident with a drunken driver and is killed. ²⁴

Fischer et al claim that this qualifies as a case of harmful direct agency since harm does come to Sam as a result of Jack's deliberately involving him in something in order to further his own purpose. But, they claim, Sam is not "the sort of agent to which the DDE is allegedly applicable". DDE (and presumably DDE*) is usually taken to apply to pairs of cases. But let's put this aside. We can still understand Fischer et al's claim that this case causes trouble for DDE*. But this case exploits Quinn's failure to make explicit that DDE* applies only to cases of foreseen harm. Jack does not foresee that Sam will be harmed on his way to the coffee-house. On my interpretation of DDE*, what Jack does fails to satisfy the conditions for either harmful direct agency or harmful indirect agency since the case fails to meet condition (i') (p. 65)

²⁴Fischer et al, p. 712

That is not to say Jack does nothing wrong, but simply that DDE* does not apply.

Fischer et al acknowledge that Quinn probably did intend for DDE* to be restricted to foreseen harms. They write:

It may seem that it would be more charitable to amend Quinn's formulation in light of what he says at the beginning of his paper, in footnote 3. There he says that the doctrine of double effect states necessary conditions of the permissibility of doing something when one foresees a bad upshot. Indeed it is reasonable to think that Quinn wishes his explicit formulation, which he develops in the text of his paper, to be embedded within a set of presuppositions which includes the constraint that the agent foresees the relevant bad upshot.²⁵

Fischer et al go on to claim that even if we interpret Quinn in this more charitable way (were we assume DDE* applies only to cases involving foreseen harm), DDE* is implausible. Here is an example of the second type of objection they raise. We are to assume (from a previous case) that Jack (not Bill) has sent Sam to the coffee house.

Bill can be taken to foresee with high probability that Sam will be killed. He expects Sam to be killed in an accident with a drunk driver. But just in case Sam avoids this fate, Bill drives to the parking lot [of the coffee house] with the intention of killing Sam, if Sam miraculously avoids the drunken drivers. Suppose further that just before Bill would have been called upon to finish him off, Sam is distracted by Bill's flashy car and [swerves over some railroad tracks where he is run over by a train and dies.]²⁶

²⁵Fischer et al, p. 715

²⁶Fischer et al, p. 716 (I have altered the case in minor ways in the interest of simplicity.)

Of this case they write: "Bill's agency is intuitively a paradigm case of direct agency: he is tugged and guided by evil... And yet Quinn's account must classify Bill's agency as indirect."

On my understanding of HDA, this is not harmful agency at all--it's neither direct nor indirect harmful agency. Here too condition (i') simply is not met. Bill does not harm Sam.²⁷

We come now to the third sort of objection. Fischer et al claim that the previous objections were meant to show DDE* is descriptively inadequate. Now they offer two objections meant to show DDE* to be normatively inadequate; I will consider both of these objections. The objections revolve around cases of bomb removal. Each of the objections involves a pair of cases in which either DDE* makes a normative discrimination that intuition doesn't make. Here is a simplified version of their first case:

BOMB REMOVER 1: On a table in the library are two boxes: a red box containing valuable books and a black box containing a bomb. In order to save the books, we ask Mary, who knows nothing of the situation, to take a steel case to the library and to put the red box (containing the books) into the steel case. We foresee the likelihood of Mary's death since we know that the bomb is sensitive to movement. Mary complies; the books are saved and she is killed in the explosion.²⁸

According to Fischer et al, this is a case of harmful direct agency. They write: "We foresee a harm to Mary, and the harm comes about. It comes about partly as a result of our involving her in removing the black box in order to further our purpose of saving the books..." Now here is the second case:

²⁷Notice that even if one opts for an interpretation of harmful direct and indirect agency that substitutes something weaker for my (i'), this case is unlikely to count as harmful agency at all.

²⁸This is a simplified version of Fischer et al's Bomb Remover 1, p. 717. The changes I have made do not affect my point.

BOMB REMOVER 2: On a table in the library are two boxes: a red box containing valuable books and a black box containing a bomb. In order to save the books, we ask Mary, who knows nothing of the situation, to take a steel case to the library and to put the red box (containing the books) into the steel case. We foresee the likelihood of Mary's death since we know that the bomb is sensitive to movement. On her way to the library, Mary forgets what she is supposed to do. She picks up the black box and carries it away with her. The bomb goes off, killing her. The books are saved.²⁹

Fischer et al claim that this is a case of indirect agency, quoting Quinn's original definition. They write: "we foresee a harm to Mary, and the harm comes about. We intend to involve her in something 'in order to further [our] purpose precisely by way of [her] being so involved.' But 'what is so intended does not contribute to [her] harm,' for what is so intended is that she put the [red] box in the case and that does not happen. Since it does not happen it cannot play a causal role in contributing to her harm." They note that while one might claim that since Mary does do something which prevents damage to the vermilion box, perhaps this is close enough to count as what was intended for her. But this puts us 'back in the game of distinguishing intentions and ascertaining "closeness". Notice though that the issue of closeness in a different place here. In CITY BOMBER, the issue of closeness arose with regard to whether the bomber intended to involve the civilians. Now it arises with regard to whether what is intended is what contributed to harm. (On my interpretation of Quinn, CITY BOMBER raises an issue of closeness around condition (ii), whereas this case raises the issue for condition (iii) (see page 65)).

²⁹This is a simplified version of Fischer et al's Bomb Remover 2 p. 718. The changes I have made do not affect my point.

According to DDE* (part-one), the first case of agency is morally worse and thus harder to justify than the second.³⁰ Fischer et al write: "Apparently, on Quinn's account [Mary's] mistake makes our action more easily justifiable, which seems implausible." I agree that this seems implausible. Further, I agree that this case is troublesome for DDE*. What should be said in defense of DDE* is that like cases can be created to cause trouble for DDE. Wherever the problem for DDE* resolves around what exactly is intended, one can concoct parallel problems for DDE.

The final pair of problem cases contrasts BOMB REMOVER 1 with the following case:

BOMB REMOVER 3: At one end of the library is a table at which Mary works. At the other end of the library is a table on which are scattered valuable books. Also on the table with the books is a black box containing a bomb. The bomb is on a timer and will blow up, we realize, in a matter of minutes. We grab the black box (containing the bomb) and carry it quickly to Mary's side of the library, where we gingerly place it on the table at which Mary is working. We run from the scene, leaving Mary to her fate and realizing that there is an extremely high probability of her death. The books are saved. Mary is killed when the bomb blows up.³¹

Fischer et al compare this case and BOMB REMOVER1 in the last of their objections I will consider. They claim that BOMB REMOVER3 is a case of harmful direct agency, while BOMB REMOVER1 is a case of harmful indirect agency. Yet in each case the mix of good and bad effects is the same; in each case, Mary's is killed and in each case the books are saved. Quinn is

³⁰DDE is usually qualified to apply only to cases where the harm brought about as a means is less than (or perhaps equal to) the good brought about as an end. DDE* might be so modified. I discuss this further after introducing the final objection.

³¹This is a slightly altered version of Fischer et al's Bomb Remover 3, p. 719.

apparently committed to the view that the what is done in BOMB REMOVER3 is more easily justified and this is implausible.

The first thing to note is that application of DDE (the original doctrine) is usually restricted to cases in which the harm brought about as a means or side effect is outweighed ,or at the very least matched, by the good end, so DDE would not normally be applied here at all. Second, BOMB REMOVER3 does not point to a special problem with DDE* as opposed to DDE. Rather, it raises a question about how best to formulate the doctrine (whether the original version or Quinn's). . Quinn could add the same sort of qualification to DDE*. However, there are two ways one might qualify the doctrine to account for proportionality and only one will help. Let's look at the two ways of qualifying the doctrine.

One might hold the following view:" For any two cases which are alike except that one involves direct agency and the other indirect, the first case is morally worse. But (here's the qualification) neither case is permissible unless the relevant proportionality requirement is met." Notice that this still commits one to the view that the agency in BOMB REMOVER 1 is morally worse than that in BOMB REMOVER 3, which I think is an undesirable commitment.

The second way to include the qualification is to state DDE* (or DDE) as follows: "It is sometimes permissible to bring about a harm by indirect agency which it would not be permissible to bring about by direct agency." The "sometimes" would then get filled out to include the proportionality requirement. This formulation does not commit one to the view that the agency in BOMB REMOVER 1 is morally worse than that in BOMB REMOVER 3. It leaves open the view that direct agency makes for moral

worseness only if the proportionality requirement (and perhaps other conditions) are met first.³²

If DDE is stated so as to entail that for any two cases alike except that one involves an intended harm and the other a merely foreseen harm, the former is morally worse, than I believe objections similar to the bomb cases can be found. On the other hand, DDE and DDE* can avoid such objections depending upon how they are qualified.

8. Summary

We have been looking at Quinn's replacement principle for DDE because it promised to do the normative work of DDE while escaping the serious problem raised by Bennett which I discuss in Chapter 2. Recall that the replacement principle was supposed to get around Bennett's objection and do a better job than the original principle at explaining certain asymmetries. I think DDE* does not correctly explain the asymmetries, since I think opportunistic agency is not morally worse than eliminative agency. So I think we should reject part-two of DDE*. I have left open whether or not DDE* (part-one) might have advantages over DDE. I have argued that as long as we interpret Quinn charitably, we need not see DDE* as being subject to the criticisms made by Fischer et al. On the other hand, DDE* does not have a couple of the key advantages Quinn suggested it to have. Most importantly for my purposes it does not get around Bennett's objection. So I will return in the next chapter to DDE and see if it can't be defended against Bennett's challenge after all.

³²Presumably one could qualify the doctrine in this way without bringing in permissibility. On this formulation the moral worseness of intending harm over foreseeing it holds only once the proportionality requirement is met. But I don't know of anyone who in fact defends or discusses this formulation.

Chapter Four: Refutation of Bennett's Challenge

1. Introduction

DDE* was of interest because of Quinn's claim that it got around Bennett's objection. In the previous chapter we saw that it in fact does not escape Bennett's objection. In this chapter I take another look at Bennett's objection and argue that we can block his *reductio* after all.

In Chapter 2, we looked at an argument that went as follows:

- (1) If feature F is a conjunction of features G and H, you don't intend to produce an F unless you intend to produce a G and intend to produce an H.
- (2') Killing the civilians involves making their bodies inoperative for a while and making their bodies inoperative thereafter.
- (3') The terror bomber does not intend to produce an event with the feature *civilians inoperative after the war*.
- (C') Therefore, the terror bomber does not intend to kill the civilians.

Let's begin by reexamining (3'). Why think (3') is true? In chapter 2, I claimed that (3') of Bennett's argument relies upon the counterfactual test question. Bennett assumes that the terror bomber does not intend to produce civilian inoperativeness after the war since the bomber would have been no less motivated to bomb had he believed the civilians would not remain

inoperative. So Bennett assumes that an agent who answers the question "If you had believed there would be no G as a result of your A-ing would you have been less likely to A?" negatively does not intend to bring about a G. I think Bennett is wrong to make this assumption. In the following section I will argue that this assumption is wrong because it does not take into account the possibility of 'preemption'.

2. The Problem of Preemption

Bennett assumes that a negative answer to the question shows that the agent is not motivated by the expectation of a G and thus does not intend to produce a G. But this does not follow. Bennett's counterfactual test question does not allow us to distinguish whether the agent answers the question negatively because she is unmotivated by the expectation of a G or whether the agent answers the question negatively because although she is motivated by the expectation of a G, she thinks in the absence of a G something else will bring about her end. The difference is important. I will be arguing that in the first case the agent does not intend to bring about a G as a means to her end and that in the second case she does intend to bring about a G as a means to her end, and that the counterfactual test does not make this important distinction. First though I want to offer an example which should help to illustrate the problem with the test question. Consider the following:

HOUSEBUILDER: Marge has decided to please Frank by building him a house. Frank is away for six months and Marge, being an exceptionally efficient carpenter, is certain she can finish the framing and the facades before he returns if she works on it every day. She believes that when Frank returns and sees the

completed facades of the house, he will be pleased. She also thinks that if she were to finish only the framing during Frank's absence, so that upon his return he would see the framing but not the facades, he would be equally pleased. Finally, she realizes that if she were not to have even the framing done, then Frank would not be pleased. So if she thought that she couldn't complete the framing in a six-month period, she wouldn't undertake the project at all. Still, she believes that she will be able to complete not only the framing but the facades as well by the time Frank returns.

Here's why this case makes trouble for Bennett's test. If Marge had believed that she would not be able to finish the facades, she would not have been less likely to undertake the project. But this, according to Bennett, means that she doesn't intend to finish the facades. I think it is clear that Marge does intend to finish the facades. She plans to finish both the framing *and* the facades by the time Frank returns. If someone asks Marge why she is working so hard on the house, she will respond "in order to have the framing and facades finished by the time Frank returns."

According to Bennett's view, the relevant question in this case is: "If you believed you would not finish the facades by the time Frank returns, would you be less likely to go ahead?" Assuming Marge answers No, we will not be able to distinguish whether she answers No to the question because (1) she thinks that the finished facades play no role in bringing about Frank's happiness (they are merely a side effect of her activity) or (2) although she thinks the finished facades will play a role in bringing about her end, she also

thinks that in their absence, something else would bring about Frank's happiness.

Consider line 1 in the diagram below. Line 1 is meant to illustrate a case in which an agent expects an action of type A to produce events of type F and G, which in turn are expected to produce an event of type H. The 'counterfactual' column indicates that the agent further believes that if no event of type G were to occur, her end, an event of type E, would still be brought about via an H. Notice that in line 1, the E is shown as coming about via an H in both columns; that is, whether or not a G occurs, an H occurs and brings about an E.

	<u>actual (a G occurs)</u>	<u>counterfactual (no G)</u>
line 1	A -> F&G -> H -> E	A -> F -> H -> E
line 2	A -> F&G -> H -> E	A -> F -> J -> E

Now turn to line 2. It is just like line 1 under the actual column but differs in the counterfactual column. Line 2 shows that if no G were to occur, an E would be brought about by a J, not an H. This is supposed to represent preemption. Whether or not a J occurs depends upon whether a G occurs. In the absence of a G, a J occurs. But if a G occurs, it preempts the J. (In this case the dependence is asymmetrical. Whether a G occurs does not depend on whether a J does). In HOUSEBUILDER, finishing the facades preempts Frank's seeing the framing. So HOUSEBUILDER is best represented by line 2 while STRATEGIC BOMBER is best represented by line 1. The strategic bomber thinks that whether his raid results in any deaths or not, it is the shortage of weapons which will bring down enemy morale.

My view is that in line 2 sort of case, the agent intends to bring about an F and G as a means, whereas in a line 1 sort of case, the agent does not

intend to bring about a G as a means. I support this by pointing out that only in line 1 does the agent expect her end to be brought about in the same way, whether or not a G occurs. The route to an E shown under the counterfactual column of line 2 involves a potential alternate means (a J); not an actual means. In HOUSEBUILDER, Marge would intend to bring about an E via a J if her beliefs were different than they actually are. But as things are, she does not intend to bring about an E via a J. But others may disagree with me. They might point out that in both cases the agent takes the production of a G to be incidental to the achievement of her end. After all, the agent thinks that her end would be achieved by A-ing, whether or not a G occurs. So the agent thinks that achieving her end does not depend upon producing a G. And surely this provides strong reason to see the production of the G as unintended, as merely foreseen. I call their view 'the Ahistorical View'. This view says:

Ahistorical View: If an agent is no less motivated to A given the counterfactual belief that A-ing produces no G, then she does not intend to produce a G by A-ing.

The Ahistorical View is implausible; it is implausible because it assigns to agents intentions which conflict with beliefs. HOUSEBUILDER provides an example of the problem. To see the point we need to consider the case in a little more detail.

Potential Alternate Means-End Route

work every day for 6 mo. → finish framing → Frank sees framing → Frank is happy

Actual Means-End Route

work every day for 6 mo. → finish facades → Frank sees facades → Frank is happy

Marge has the following beliefs (as suggested by the diagram): She believes that if she finishes the framing and facades, then it Frank's seeing the facades which will make Frank happy. Furthermore she believes that she will finish the facades. Now recall that according to the counterfactual test, Marge intends to complete only the framing. So if we accept that Marge intends to finish the framing only as a means to making Frank happy, then we are left assuming Marge (1) intends to make Frank happy by finishing the framing and not by finishing the facades and (2) believes that he will not see the framing, but will see the facades. Thus it leaves her intending to bring about her end by bringing about a means (in this case Frank's seeing the framing) which she does not expect to occur. I think this gives us a strong reason to reject the Ahistorical View.

Philosophers often move from claims about what an agent needs (in order to bring about her end) to what the agent intends. For example, Warren Quinn, summarizing Bennett's point, writes: "The terror bomber does not...need the civilians actually to be dead. He only needs them to be as good as dead and to seem dead until the war ends."¹ But it should now be clear that such appeals are misleading. *It does not follow from an agent's not needing a G in order to bring about her end that she does not intend to produce a G as her means to that end.* Again, Marge does not need to finish the facades in order to achieve her end, but it does not follow that she does

¹Quinn, "Actions, Intentions and Consequences: The Doctrine of Double Effect", 337.

not intend to finish them as her means. And more importantly, even if the terror bomber does not need civilians to die in order to achieve his end, he may still intend to kill as a means.

I have been arguing that there is a problem with Bennett's counterfactual test question. How does this bear on Bennett's reductio? The counterfactual test was supposed to support (3') of Bennett's argument. It was supposed to support the claim that the terror bomber does not intend to make civilians inoperative after the war. If I am right that there is a problem with the counterfactual test question, then this undercuts one line of support for (3'). But notice it does not show (3') is false. (3') may be supportable on other grounds. One might defend (3') by appeal to the following plausible principle:

If an agent expects a G to occur only after her end, an E, then she does not intend to bring about a G as a means to an E.

Since the terror bomber expects the civilian inoperativeness after the war to occur after the terror, he does not intend to bring it about as a means to the terror.

While it is important to notice the problem with Bennett's counterfactual test, the availability of alternate defenses of (3') shows that this test is not what is doing the work in Bennett's argument. We have been assuming that (1) and (2') entail the following intermediate conclusion: If one does not intend to make the civilians inoperative thereafter, one does not intend to kill them. In chapter 2, I mentioned that Bennett defends his argument against the response which simply rejects temporal disjunction. Bennett's reply was that if we disallow temporal disjunction we seem to block the reasonable conclusion that the politician (in Bennett's example) intends

to break up the union for only the month of December. Recall that Bennett's case, which I will henceforth call 'STRIKE' was this:

STRIKE: A political leader takes action against a trade union, intending to bring about a month-long state of disintegration in which the various locals break off from the parent body and severally fall into further disunity. This is his intended means to the end of the union's being unable to call a strike during December. He is rightly sure that if the union falls apart for that long it will never be reconstituted, but all he cares about or intends is the one-month dissolution: if he were sure that the union would recover during January and flourish for many years, that would not reduce his motivation for moving against it.

Bennett argues that it is reasonable to think that the politician intends to bring about only a one-month break up of the union. He uses this claim to defend temporal disjunction. As I noted in Chapter 2, this sort of move is suspicious even in cases which do not involve irreversible change. Do I intend to get four hours of cold relief only if I intend to get the first hour of relief and I intend to get the latter three hours of relief from my cold? Though I reserve the name 'temporal disjunction' for cases involving permanent effects (as above), this whole sort of strategy is suspicious and it would be nice to reject it while maintaining that the politician intends to bring about only a one-month dissolution. I think there are grounds for doing so, and I turn to those in the following section.


3. A Dilemma

I believe that Bennett's defense of temporal disjunction relies upon an unclarity in the structure of the relevant cases. Bennett relies heavily on the STRIKE to motivate his argument. Presumably we are to assume that STRIKE is similar to TERROR BOMBER in all the relevant ways. But on the most natural reading of STRIKE, the case's structure differs importantly from the structure of TERROR BOMBER. Consider these lines:

[The politician] is rightly sure that if the union falls apart for that long [one month] it will never be reconstituted, but all he cares about is the one-month dissolution....


This suggests that the politician thinks the one-month period of inoperativeness brings about the inoperativeness thereafter. It implies that if the union were only disassembled for a day or a week, it would be reconstituted. The politician also thinks it's the one-month break up that prevents a strike in December. So he thinks of the one-month break-up as playing two causal roles: bringing about a permanent dissolution of the union and preventing the strike. The diagram below depicts this way of understanding the case. The diagram shows the strike's being averted and the union being dissolved from January 1 on as flowing from the one-month dissolution.

Diagram #1



What about TERROR BOMBER? If STRIKE is to be of any help to Bennett, then it and TERROR BOMBER need to be alike in all the relevant ways. The relevant bomber case is one something like what is depicted in Diagram #2.

Diagram #2




I believe Bennett faces a dilemma. He can say that Diagram #1 and Diagram #2 do depict the cases he had in mind or he can deny that the natural reading of STRIKE is the one he had in mind and claim that the two cases are structurally alike but not as in the diagrams above. I will be arguing that Bennett runs into trouble no matter which of these he chooses.

First assume that the cases are as shown in the diagrams above. I said that a natural reading of STRIKE suggests such a case. But what about TERROR BOMBER? For a terror bomber case to have this structure we have to assume that the bomber expects to produce civilian bodily inoperativeness which is not yet permanent (so not yet death). Why? The condition brought upon the union (the inoperativeness for December) is a reversible condition. The union is disassembled but it is physically possible for the union to reassemble. So, for the cases to be analogous, the bomber needs to assume that the effect produced on the civilians is reversible. Thus we need to assume

that the bomb produces inoperativeness which is not death and further, that the civilians' remaining in this reversible state of inoperativeness for a certain length of time will bring on an irreversible state of inoperativeness (or death). We have to assume all this in order for the cases to be parallel (given that Diagram #1 accurately depicts STRIKE).²

But an ordinary case of terror bombing is not like this. Most notably, it is not usually assumed that the bomb brings about inoperativeness which is not yet death. We need a description of a case which truly parallels STRIKE (and goes with Diagram #2). I happen to have such a case, but before turning to it, I need to say a little about stun bombs. Stun bombs, we are to assume, are a special kind of bomb. People hit by stun bombs typically go into a state of temporary inoperativeness (they become stunned). The inoperativeness lasts for a few days. If at this point the victims are treated, they fully recover. If they are not treated, they die as a result of the inoperativeness. Now consider the following case:

²Notice that the problem of preemption arises for the Counterfactual Test Question in cases in which we are to assume that both effects (the dying and the factory destruction from example) flow from the agent's action (only in cases which have the second structure). It won't arise if the cases are as shown in Diagram #1 and #2. It does not arise for cases with the structure shown under the actual column below:



Recall that the problem with the test was that it could not separate when an agent was actually unmotivated by the belief that A-ing would lead to a G and when she was motivated by this but believed that in the absence of a G an alternate route from the F (see diagram) to her end would occur. So the trouble comes about when the agent believes that in the presence of a G, an F will bring about an E by one route and in the absence of a G, an F will bring about an E via an alternate route. In the above diagrammed case, the second possibility (the alternate route) is ruled out. The reason is because here an F does occur without a G and it is stipulated then when this happens, the F brings about a G.

STUN BOMBER: A bomber sees that the enemy is on the brink of surrendering; he aims to push them over the edge by targeting civilians in a small village. His plan is to drop stun bombs on the inhabitants of the village. He is sure that the inoperative bodies will be found within a day or so and that the enemy will surrender almost immediately. He would be happy for the stunned civilians to recover after the surrender. But he also knows that the enemy is ignorant about stun bombs. So he knows that the civilians will go untreated and thus will die.

Now we have a case which is parallel in structure to STRIKE. So now Bennett's point is in order: if we want to maintain that the politician intends to bring about only a one month break up, we will have to say the same of the stun bomber--he intends to stun the civilians, but not to kill them. (Otherwise put, he intends to make them inoperative for a few days). But it seems to me that this is the right thing to say in the case of the stun bomber--or at least it is not an absurd thing to say. Stun bombers are not terror bombers. It is the case of the terror bomber which I said offers the paradigm of intending to kill as a means. But the stun bomber case is importantly different from an ordinary case of terror bombing. In making the stun bomber case parallel to STRIKE we had to assume that the bomb causes inoperativeness which is not yet death. This assumption is not present in a standard account of terror bombing. The stun bomber's intending only temporary inoperativeness does not make trouble for the notion 'intended as a means'. If Bennett really has a stun bomber in mind, his reductio has lost its force. That's the first horn of the dilemma, now let's turn to the second horn.

Bennett might deny that he has cases in mind like those in Diagram #1 and Diagram #2. He might claim that the relevant cases are rather as shown below.

Diagram #3




Diagram #4




Diagram #3 differs from Diagram #1 in that the effect *union dissolved from Jan. 1 on* follows directly from the politician's action rather than from the one month dissolution. Likewise, Diagram #4 differs from Diagram #2 in that the civilian inoperativeness from time t on follows directly from the bomber's action and not from the inoperativeness up to time t.

Given that the cases are structured as in Diagram #3 and Diagram #4, I think it is again open to us once again to reject temporal disjunction. The terror bomber expects his action to result in civilians dying and the politician expects his action to result in the union dying. In each case, the agent believes that something less would do to bring about their end (inoperativeness for a

while), but that does not show they are not motivated by the expectation of the relevant death or deaths. Bennett thinks we should be reluctant to attribute to the politician the intention to kill the union as a means. But if Bennett means for STRIKE to have be like the Diagram #3 case, then I think any reluctance we feel to attribute to the politician an intention to kill the union stems from Bennett's misleading description of the case. As I have previously noted, Bennett's description of STRIKE strongly suggests a case more like the one shown in Diagram #1 than the one shown in Diagram #3.

I am rejecting (2') of Bennett's argument but remember that we are giving Bennett (3'). (3') alone gives Bennett the conclusion that the terror bomber does not intend to make the civilians inoperative after the war. But without temporal disjunction, this claim is compatible with the terror bomber intending to kill the civilians. Likewise, that the politician does not intend anything for the union in January does not show he does not intend to kill the union as his means. This case seems to be one of a group of odd cases. These odd cases involve agents adopting means to their ends which are 'too strong'. Here's an example:

It's very important to you that you sleep for two hours but you are not at all tired. The only way you can see of getting to sleep is to take a sleeping pill and the only sleeping pill you have available to you is a four hour sleeping pill. You take the four hour sleeping pill as your means to sleeping for two hours.

Do you intend to sleep for hours three and four? Your belief that you will do so as a result of taking the pill is certainly not what is motivating you. Whether or not you do, what is clear is this: you do intend to take a four hour

sleeping pill. You would take a two hour pill if you had one, but you don't. So you intend to take a four hour pill. This case is not importantly different from TERROR BOMBER. The terror bomber would drop a stun bomb if he had one, and bring about terror by stunning. He would bring about terror by doing something which is reversible (bringing about mere inoperativeness), but he believes that the only way he can bring about terror (and bring about inoperative bodies) is by producing an irreversible effect on the civilians. So he plans to drop a conventional bomb instead. He intends to kill as his means to creating enemy terror. That the terror bomber does not intend to make anything the case for the civilians after the war does not show he does not intend to kill the civilians as his means. To claim otherwise would be like claiming in the case above that you don't really intend to take a four hour pill, or like claiming that you don't intend to eat an apple from the bowl unless you intend to make it the case that the apple not be in the bowl tomorrow (and that it not be in the bowl the day after, and so on).

That's the second horn. Let me summarize. If Bennett claims that the cases have the structure shown in Diagrams #1 and #2, then he seems to have in mind a stun bomber, not a terror bomber and the reductio loses its force. If Bennett claims that both cases have the structure shown in Diagrams # 3 and #4, then it is plausible to think that the politician really does intend to kill the union as his means and that any intuitions to the contrary are due to Bennett's misleading description of the case.

4. Summary

My purpose in this chapter has been to argue that we can reject Bennett's reductio. I have suggested that we reject it by rejecting the move to resolve features of the form 'F forever' into ones of the form 'F for a while'

and 'F thereafter'. Furthermore, I have stressed that we should also reject his counterfactual test question, and rough appeals to what an agent needs, as indicators of what agents intend as means. Such tests cannot account for preemption and thus will sometimes indicate not what an agent does intend but what an agent would intend.

So far I have focused on rejecting challenges to the Ordinary View. I have not offered any sort of positive account of 'intending as a means'. While I will not discuss a positive account in any detail, I will briefly sketch some ideas in this direction in the following chapter.

Chapter Five Intending as a Means

1. Introduction

So far I have claimed that what any test for what agents intend as means needs to pick out agents' actual means, not potential alternate means. In many respects, I think Bratman's view of intending is satisfactory and I will suggest below how with the rejection of Bennett's challenge, it may be defended from the objections I raised earlier.

Bratman relies upon counterfactuals much like the one Bennett offers in explicating the endeavoring and screening-off roles he claims are characteristic of intentions. I have said that Bennett's counterfactual test should be rejected. Where does that leave Bratman's account, or other accounts which identify intentions mainly with what Bratman considers endeavoring? I will address this question below.

2. Means Consistency Principle

The following principle offers a plausible constraint on what is intended as a means:

Means Consistency Principle: If X believes that by A-ing she will not produce an N, then she does not intend to bring about an N as a means to her end for which she is A-ing.

I take it that this principle is uncontroversial. This principle will lead us away from the Ahistorical View. Bennett's counterfactual test gave results which conflicted with this principle. But perhaps a revised version, like the one below, would not.

3. Revised Counterfactual Test

I have been arguing that in line 1 (see below), the agent does not intend to bring about a G as a means to an E; the G is a side effect of A-ing and in line 2 the agent intends as her means an F&G. The problem with Bennett's test question ("If you had believed that there would not be a G would you have been less likely to A?") was that it could not pull these two types of cases apart.

	<u>actual (a G occurs)</u>	<u>counterfactual (no G)</u>
line 1	A -> F&G -> H -> E	A -> F -> H -> E
line 2	A -> F&G -> H -> E	A -> F -> J -> E

But consider a revised test which will separate line 1 type cases from line 2 type cases. The revised test involves appealing to two counterfactual questions instead of one; the idea is that the test can be applied for any case with two effects flowing from the action where the question is whether one of these is intended. The questions (which refer to the chart above with line 1 and line 2) below provide a test for whether the agent intends to bring about a G.

- (Q1) If you had believed that as a result of your A-ing there would be an F and a G but no H, would you have been less likely to A?
- (Q2) If you had believed that as a result of your A-ing there would be an F and a G but no H and no G, would you have been less likely to A?

What will be important is whether the agent's responses to the two questions differ. If the agent answers Yes to Q1 and No to Q2, this suggests that there is preemption involved. However if the agent answers Yes to both questions then this shows that she intends to bring about a G as a means (or part of a means). If an agent answers No to both questions then she does not intend to bring about a G as a means. This test should work for any cases which has the structure shown under the actual column above. To make this clearer, I will return to HOUSEBUILDER and show how this strategy works where Bennett's question did not.

Earlier I diagrammed HOUSEBUILDER as follows:

work every day for 6 mo. → finish framing and facades → Frank sees facades → Frank is happy

Marge's activity results in a finished frame and finished facades. But recall that the case is like line 2 not line 1 because Marge supposes that if she were not to finish the facades her end (Frank's being pleased) would come about in a different way (via Frank's seeing the frame). And the problem I raised earlier was that Bennett's test question could not distinguish whether HOUSEBUILDER was like line 1 or line 2.

The revised test involves comparing Marge's answers to the following two questions:

- (Q1) If you had believed that Frank would not see the facades, would you have been less likely to undertake the project?
- (Q2) If you had believed that Frank would not see the facades and that the facades would not be done, would you have been less likely to undertake the project?

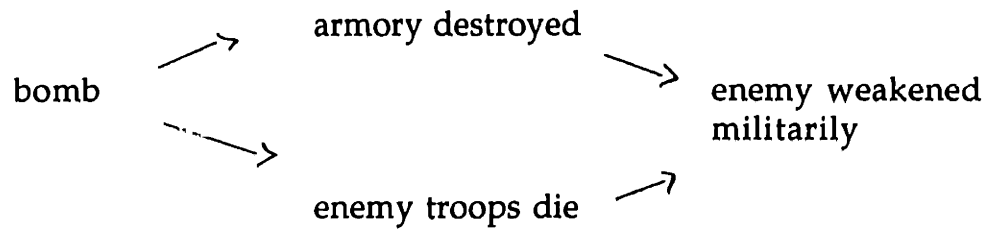
Here the answer to (Q1) is Yes (since, remember, she believes the facades will be completed and thus that it is seeing the facades which will make him happy). And her answer to (Q2) is No. What's important is that her answers to the two questions differ. If she had believed not only that he wouldn't have seen the facades but additionally that she wouldn't have finished them, then she would not have been less likely to work on the house, since this leaves her believing that she finishes the frame and he sees the frame. Again what is important is not her answer to either question in isolation but the difference in the responses. The difference brings out that she thinks in the absence of finishing the facades, her end would be brought about by an alternate causal route. If she did not believe that finishing the facades preempted an alternate route, then she would give the same response to each of the questions.

The revised test lacks simplicity; it's a burdensome test to use. Still, it, or some other revised version of the counterfactual test strategy may succeed at handling preemption.

4. Means-End Chains

The view I am pointing towards has agents intending each link in what they take to be a causal chain from their means to their ends. It is tempting simply to assert the following: *If an agent believes that by A-ing she will produce an F and she believes that the F will help to bring about her end, E, then she intends to bring about an F as a means.*¹ "Help to bring about" is undesirably vague, but there is a more immediate problem with the view. This is that it does not allow for cases where there are two separate routes to an agent's end. For example, so far we have assumed that the deaths the strategic bomber foresees he'll cause do not contribute to his end, only the weakening of the armory contributes to his end. But now let's assume that he sees he'll cause not civilian deaths, but deaths of enemy troops. He also sees that this will help to accomplish his end of weakening the enemy. So by dropping his bomb he believes he will bring about destruction of the armory and the deaths of enemy troops and he believes both of these effects will contribute to the weakening of the enemy. The case might be diagrammed as follows:

¹This turns out to be quite close to the view Bratman advocates.



Now assume additionally that the bomber intends to weaken the enemy by destroying the armory. He foresees that enemy troops will die and he knows that this effect will in fact contribute to his end, but it is not his chosen means to his end. Furthermore, assume he believes that if the troops were not to die, the armory destruction would weaken the enemy in the same way as if they do die. So what we have is a case of two independent causal routes to the bomber's end. It does seem to me quite plausible that the bomber intends only the armory destruction as his means. That is why the view (in italics above) won't do. Bratman suggests that agent's intend *at least one* means-end chain. (For example, in the case above, if the bomber intends to produce his end by bombing, he intends either the chain which brings about his end via the deaths or the one that goes via the armory destruction, or both; but he must intend at least one of these.)² I think this is right, but let's consider one final problem.

Consider a case which involves what might be called 'tracking'. Imagine a slight variant on the strategic bomber case. The raid is to occur at night. In order to find his target, the bomber will look for a hospital which is right next to the armory. Why will he look for the hospital? Because it is well lit at night. He will endeavor to drop his bomb over the lights of the hospital.

²Presumably he might also intend that his end come up in one way or the other.

He will reason about finding the lights, since they are his only clue to finding the armory. He may even use the destruction of the hospital as an indicator as to whether he's been successful at hitting the armory. If, for example, after dropping a bomb, he notices the lights are still on, he may fly back and drop more bombs. Does he intend to destroy the hospital? On Bratman's account it appears that he does. I'm not sure this is a problem for Bratman's account. It's difficult to know what to say about such a case. Perhaps in such a case the bomber does intend to hit the hospital, or to knock out the lights. But it is worth noticing that he does not intend to do so as an end; his end is to hit the armory. And, he does not intend to do so as a means. He does not believe that the destruction of the hospital causes the destruction of the armory. He is simply using the hospital's destruction as a guide to destroying the armory. Hitting the hospital lights seems to be something like an epistemological means; in any case, it's not a causal means. But perhaps some effects are intended neither as (causal) means nor ends.

Conclusion

My purpose in the thesis has been to argue against two challenges to the Ordinary View, and thus to the Doctrine of Double Effect. The holistic challenge put pressure on the scope of intentions from one direction--it suggested that the scope of decisions, and likewise intentions, is broader than we might ordinarily think. Bennett's challenge to the notion "intending as a means" relied on pressure from the other direction on the scope of intentions. He relied upon the claim that in delimiting the scope of intentions, we must "shave things as closely as possible". I have argued that both of these pressures can be blocked. We can block the pressure from the holist camp by rejecting that decisions themselves are holistic. And we can reject Bennett's challenge, in part because he relies on an overstrong test for what is intended (one that picks up what an agent would intend), and in part because he relies upon a faulty analogy between two key cases (TERROR BOMBER and STRIKE).

I have also considered a replacement principle for DDE offered by Warren Quinn as an attempt to get around Bennett's objection. I have claimed that the replacement principle is also subject to Bennett's objection so does not do what Quinn hoped it would do. However, I think the replacement principle offers an interesting alternative to DDE. To decide

between them, one would need to take into account what sort of rationale could be given for each, and that is something I have not addressed.

I have said very little by way of a positive account of what it is to intend to bring about some effect as a means, versus to merely foresee it, (or to see it as a side effect). Nor have I attempted to defend the Doctrine of Double Effect itself. My aim has been simply to show that one kind of attack on the Doctrine of Double Effect, namely the kind that impugns the distinction between what we intend and what we foresee, fails.

Bibliography

- Anscombe, G. E. M., *Intention*, 2nd edn, Basil Blackwell, Oxford, 1979.
- Anscombe, G. E. M., "War and Murder", in *Ethics, Religion and Politics*, University of Minnesota Press, Minneapolis, 1981.
- Audi, R., "Intending," *Journal of Philosophy*, 70 (1973), 387-402.
- Bennett, Jonathan, "Morality and Consequences", *Tanner Lectures on Human Values*, University of Utah Press, Salt Lake City (1981).
- Bennett, Jonathan, "Whatever the Consequences" in Fischer et al, *Ethics, Problems and Principles*, HBJ, 1993.
- Bentham, Jeremy, *An Introduction to the Principles of Morals and Legislation*, 1789.
- Bratman, M., *Intention, Plans, and Practical Reason*, Harvard University Press, Cambridge, Mass., 1987.
- Bratman, M., "Davidson's Theory of Intention", in B. Vermazaen and M.B. Hintikka, eds, *Essays on Davidson. Actions and Events*, Clarendon Press, Oxford, 1985, 13-26.
- Bratman, M., "Two Faces of Intention", *Philosophical Review* 93, (1984) 375-405.
- Chisholm, R.M., "The Structure of Intention", *Journal of Philosophy*, 67 (1970), 633-47.
- Costa, Michael J., "The Trolley Problem Revisited" in Fischer et al, *Ethics, Problems and Principles*, 1992.
- Costa, Michael J., "Another Trip on the Trolley" in Fischer et al, *Ethics, Problems and Principles*, 1992.
- Danto, A.C., "Basic Actions", in A.R. White, ed., *The Philosophy of Action*, Oxford University Press, Oxford, 1968, 43-58.
- Davidson, D., *Essays on Actions and Events*, Clarendon Press, Oxford, 1984.
- Davidson, D., "Replies to Essays I-IX", in B. Vermzaen and M.B. Hintikka, eds, *Essays on Davidson. Actions and Events*, Clarendon Press, Oxford, 1985, 13-26.
- Davis, Nancy, "The Doctrine of Double Effect: Problems of Interpretation" in Fischer et al, *Ethics, Problems and Principles*, 1992.
- Davis, W.A., "A Causal Theory of Intending", *American Philosophical Quarterly*, 21 (1984), 43-54.
- Duff, Anthony, *Intentions, Agency, and Criminal Liability*, Basil Blackwell, Oxford, 1990.

- Fischer, Ravizza, and Copp, "Quinn on Double Effect: The Problem of 'Closeness'", *Ethics*, (1993), 707-725.
- Foot, Philippa "The Problem of Abortion and the Doctrine of Double Effect" in her *Virtues and Vices and Other Essays in Moral Philosophy*, University of CA Press, Berkeley, (1978).
- Frankfurt, H.G., "The Problem of Action", *American Philosophical Quarterly*, 15 (1978), 157-62.
- Frankfurt, H.G., "Freedom of the Will and the Concept of a Person", in G.Watson, ed., *Free will*, Oxford University Press, Oxford, 1982, 81-95.
- Goldman, A.I., *A Theory of Human Action*, Princeton University Press, Princeton, NJ, 1970.
- Goldman, A.I., "The Individuation of Action", *Journal of Philosophy*, 68, (1971), 761-74.
- Harman, G., *Change in View*, MIT Press, Cambridge, Mass, 1986.
- Harman, G., "Willing and Intending" in *Philosophical Grounds of Rationality*, ed R.Grandy and R.Warner, Oxford University Press, Oxford, 1986.
- Kant, Immanuel, *Groundwork of the Metaphysics of Morals.*, Translated by H.J. Paton. Harper & Row, New York, 1964.
- Lewis, David, *Philosophical Papers, Vol. II*, Oxford University Press, Oxford, 1986.
- McIntyre, Alison, "Agency, Deliberation, and Foreseen but Unintended Means" MS, (1987)
- Melden, A.I., "Willing", in *The Philosophy of Action*, Oxford University Press, Oxford, 1968, 70-78.
- Mill, John Stuart, *Utilitarianism*, 1863.
- Moore, G.E., *Ethics*, Oxford University Press, 1965.
- Moya, Carlos J., *The Philosophy of Action. An Introduction*, Polity Press, Oxford, 1990.
- Nagel, T., "Moral Luck" in *Mortal Questions*, Cambridge University Press, Cambridge, 1979.
- Nagel, T., *The View from Nowhere*, Oxford University Press, Oxford, 1986.
- Quinn, Warren, "Actions, Intentions, and Consequences: The Doctrine of Double Effect", *Philosophy and Public Affairs* 18 (1989), 334-51.
- Quinn, Warren, "Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing", *The Philosophical Review* (1989), 131-155.
- Rawls, John, *A Theory of Justice.* Harvard University Press, Cambridge, Mass, 1971.
- Scheffler, Samuel, *Consequentialism and its Critics*, Oxford University Press, Oxford, 1986.
- Searle, J. R., *Intentionality*, Cambridge University Press, Cambridge, 1983.
- Sen, Amartya and Bernard Williams, *Utilitarianism and Beyond*, 1982.
- Smart, J. J. C. and Bernard Williams. *Utilitarianism: For an Against*, Cambridge University Press, Cambridge, 1973.
- Taurek, John, "Should the Numbers Count", *Philosophy and Public Affairs*, (1977), 214.

- Thalberg, I., "Do Our Intentions Cause Our Intentional Actions?", *American Philosophical Quarterly*, 21 (1984), 249-60.
- Thomson, Judith Jarvis, "Individuating Actions", *Journal of Philosophy* 68 (1971), 771-81.
- Thomson, Judith Jarvis, *Rights, Restitution, and Risk: Essays in Moral Theory*. William Parent, ed., Harvard University Press, Mass., 1986.
- Thomson, Judith Jarvis, *The Realm of Rights*, Harvard University Press, Cambridge, Mass., 1990.
- Thomson, Judith Jarvis, "The Time of a Killing", *Journal of Philosophy* 68 (1971), 115-132.
- Walzer, Michael, *Just and Unjust Wars*, Basic Books, NY, 1977.
- Williams, Bernard, "Moral Luck" in *Moral Luck*, Cambridge University Press, Cambridge, 1981.