

WILLIAM J. FITZPATRICK

THE INTEND/FORESEE DISTINCTION AND THE PROBLEM OF “CLOSENESS”

ABSTRACT. The distinction between harm that is intended as a means or end, and harm that is merely a foreseen side-effect of one's action, is widely cited as a significant factor in a variety of ethical contexts. Many use it, for example, to distinguish terrorist acts from certain acts of war that may have similar results as side-effects. Yet Bennett and others have argued that its application is so arbitrary that if it can be used to cast certain harmful actions in a more favorable light, then it can equally be manipulated to do the same for any kind of harmful action. In response, some have tried to block such extensions of the intend/foresee distinction by rejecting its application in cases where the relation between the plainly intended means and the harm is “too close”. This move, however, has been attacked as vague and obscure, and Bennett has argued that all the plausible candidates for explicating the idea of excessive closeness ultimately fail. In this paper, I develop and defend an account of excessive closeness with the aim of rescuing the intend/foresee distinction from such charges of arbitrariness. The account is based on the distinction between *merely causal* and *constitutive* relations among *states of affairs*, and I show both how it escapes Bennett's objections to other accounts and how it applies to a variety of cases. Finally, I also examine Quinn's alternative move of shifting the focus of the intend/foresee distinction in an attempt to sidestep the issue of closeness, and argue that it is not ultimately successful. In fact, Quinn's view has shortcomings that can be resolved only by returning to an appeal to some notion of closeness, underscoring the need for the sort of account I offer.

1. INTRODUCTION

The intend/foresee distinction has long played a role in ethical thought, and while some of its traditional applications no longer have much appeal outside of the Catholic Church, it continues to be widely cited in a variety of ethical debates. Even those who reject traditional moral appeals to the intend/foresee distinction in voluntary end-of-life cases, for example, often

find it indispensable to a proper characterization of terrorism, and of what sets terrorism apart from other acts that may have similar results as unintended side-effects (Quinn, 1993b). It also figures prominently in debates over cloning for biomedical research, and not just from religious perspectives (see FitzPatrick, 2003a).

The intend/foresee distinction has typically (though not exclusively) been employed in connection with the doctrine of double effect (DDE) – the idea, broadly, that it is *sometimes* permissible to bring about *as a foreseen but unintended side-effect* some harm it would have been impermissible to *aim at as a means or as an end*, all else being equal.¹ We may distinguish, for example, between foreseen but unintended harm to civilians (“collateral damage”) caused by strategic strikes on nearby military targets, on one hand, and intended harm to civilians brought about through terrorist strikes that target civilians deliberately as a means to military, political or religious ends, on the other. Acts involving the former are certainly not always or easily justifiable, even in the context of an otherwise just war. Still, it is a common and plausible thought that such acts may at least sometimes be morally justified, while acts involving the deliberate targeting of civilians as means are either never justified or are so only in much rarer and more extreme circumstances. And while it may be possible to account for this thought in part by appeal to consequentialist considerations (“terrorism doesn’t work”), many feel a strong pull to locate the moral difference directly in the intend/foresee distinction itself – in the idea that there is something distinctly problematic about intending harm to innocents as a means to one’s ends. This seems to get at something missed not only by consequentialist accounts, but also by accounts that focus exclusively on other constraints, such as one built around the *doing/allowing* distinction, since one is plainly *doing* the harm to innocents in both bombing cases above.²

I shall not be concerned here to try to show that the intend/foresee distinction is indispensable in normative ethics, though I believe it is. Instead, I wish to examine and respond to a fundamental *challenge* that, if not successfully answered,

threatens to undermine any serious role for it. This challenge involves the charge that the application of the intend/foresee distinction is so arbitrary and flexible that if it can be used to help justify what we would like it to, then it can *equally* be manipulated to help justify almost anything. Typically, this involves arguing that if one can drive a wedge between what is intended and what is merely foreseen in one kind of case (e.g. strategic bombing of military targets), then with a little cleverness one can do the same even in what is usually offered as the contrasting case (e.g. terror bombing of civilians). The distinction therefore appears too arbitrary for principled moral application.

In response, proponents of the moral significance of the intend/foresee distinction may take either of two paths. One is to try to block such extensions of the distinction to the “wrong” cases, giving a principled account of why, for example, the deaths of civilians in the terror bombing case *cannot* properly be described as unintended in any relevant sense. Thus, one might try to show (following a suggestion from Anscombe and Foot) that the relation between the plainly intended means and the harm in such cases is simply “too close” to allow talk of aiming at the one without aiming at the other. The second option is to shift the focus of the intend/foresee distinction away from *harm*, so as to sidestep the whole issue of “closeness”. If what really matters, for example, is not whether the *harm itself* is intended as a means or end, but only – as Warren Quinn has argued – whether some *involvement* is intended as a means where that involvement is foreseen to be harmful, then we needn’t worry about clever arguments designed to show that even in the disfavored cases (such as terror bombing) “the harm itself” may not be intended (Quinn, 1993b, pp. 178–180).

With regard to the first option, many have rightly complained that more needs to be said about the obscure notion of “excessive closeness” before it can be relied upon as a principled restriction on unwelcome extensions of the intend/foresee distinction.³ Indeed, it was apparently pessimism over the prospects of such an account that largely motivated Quinn’s shift to the second alternative in his revisionary formulation of

the DDE, despite the fact that this involves a significant departure from the real focus of the DDE, as Kamm has pointed out (1992, pp. 379–380). But such pessimism is premature for two reasons. First, many cases – including the terror bombing case, which largely motivated Quinn’s revisionary move – may largely be handled just by clearing up confusions about their true structure, without even getting into the issue of closeness. (Though we will see that if one presses the case in a certain way then the issue of closeness can arise here after all and will need to be dealt with.) Second, in cases where the notion of closeness really is the primary issue, we can flesh out the notion of excessive closeness in terms of a distinction concerning the relation among relevant states of affairs, thus providing a principled way of blocking implausible applications of the intend/foresee distinction.

I shall take up these points in turn, beginning in Section 2 with an examination of how a familiar range of problem cases can largely be handled without even appealing to the notion of closeness. I will then go on in Sections 3, 4 and 5 to develop and defend a positive account of “excessive closeness” to deal with the kinds of cases that genuinely require it, beginning with relatively clear cases and then exploring a number of complications to flesh out the account. I do not pretend that the proposed account eliminates all gray areas, which I believe is a dubious aim in any case. It will be enough if it allows us to defend a wide and important range of applications of the intend/foresee distinction in ethics against the above charges of arbitrariness. Finally, I will conclude by arguing that Quinn’s alternative, while it can successfully be defended against many objections that critics such as Bennett have raised against it, has shortcomings that can be resolved only by returning to an appeal to some notion of closeness – underscoring the need for the sort of account I offer here.⁴

2. MISTAKING PROXIMATE MEANS FOR SIDE-EFFECTS

There is a class of cases that might at first seem to threaten limitless applications of the intend/foresee distinction unless we

block them by restrictions on “closeness”, but which upon examination can largely be dismissed by resolving a basic misunderstanding over their structure. A well-known example comes from Bennett (though his use of it has changed somewhat over the years): what is to stop even a terror bomber from claiming that he does not intend the deaths of the civilians he targets with lethal force, on the grounds that all he really *needs* is for them to *appear* dead for long enough to demoralize the enemy and hasten the end of the war? Since “there is nothing requiring that the people actually become dead”, why can’t the terror bomber say he *intends* only to bring about the *appearance* of death, merely *foreseeing* that his bombs will kill people? (Bennett, 1981, pp. 111–113; 1995, pp. 210–212).

In his most recent discussion of this case, Bennett calls this a “mad result” and “outright crazy” (1995, p. 211, 210), and uses the case simply to point out the failure of one particular proposal for delineating the idea of excessive closeness, which he argues would fail to block such results.⁵ Originally, however, the example was introduced as allegedly revealing a general problem for the inherently messy concept of *intending something as a means*, which “cannot be given a firm, clear theoretical grounding [that] implies what we think true and not what we think false regarding what people intend” (1981, p. 113). The suggestion was that we seem to be stuck with such absurd results due to the vagueness and arbitrariness of the concept of intending something as a means, which is precisely why he advocated jettisoning it from first-order moral thinking. And this argument has been taken seriously: as already mentioned, it is what led Quinn to abandon the traditional DDE and to propose instead a significantly revised version that he thought would avoid such problems. But is it really so hard to dispose of cases like this, so as to keep a grip on a reasonably disciplined notion of intending something as a means?

On the face of it, we can fairly easily see what is wrong with the sort of claim Bennett’s terror bomber wishes to make about his intentions, and we don’t need to appeal to any refined notion of “closeness” to explain it. The terror bomber seems just to be mistaking a more *proximate intended means* for a *foreseen*

*but unintended side-effect.*⁶ His end is to hasten the conclusion of the war, and the means he chooses to fix attention on is *to make lots of civilians on the enemy's side appear to be dead* (thus demoralizing the enemy). But the obvious question to ask is: how does the bomber intend to accomplish *that*? And the answer is: by killing them, which he does by dropping lethal bombs on them. For one very easy and effective way to make lots of people appear dead is to kill them, and that is exactly the *proximate means* the bomber has chosen to take (by the still more proximate means of dropping his bombs) to make them appear dead. Their being killed, then, is not a merely foreseen but unintended side-effect of his act, but his relatively proximate intended means to bring about the thing he prefers to focus on (their appearing to be dead), which is just a relatively less proximate means along the very same chain. The bomber drops his bombs as an intended means of killing people, which he intends as a means to make them appear dead, which he intends as a means of demoralizing the enemy, which he intends as a means of hastening a favorable end to the war.

The fact that the bomber would be just as happy if “by some miracle” the bombs he dropped only knocked people out from the shock waves (all the shrapnel having gathered, by some quantum mechanical accident, in an unoccupied corner of the village), making them appear dead for a time, is irrelevant. All that shows is that their deaths were not aimed at *as ends*, and that the bomber's end could conceivably be achieved in other ways – as with the help of such a miracle. This does not show that he was not *in fact* aiming at their deaths as *his actual chosen means* to his end. On the contrary, it is clear that he was: for having no control over miracles, and having no special “knock-out” bombs with which he could honestly try to knock people out without killing them, he simply dropped lethal bombs on a civilian population; and he obviously did this because he knew that this was a way to pursue his goals, making his targets appear dead by killing them, which is what the bombs he used do. The same may be said in the case of someone who deliberately shoots a fleeing mugger with lethal force, and then claims that his intent was only to “stop” him so

that the stolen wallet might be retrieved. It may be true that the agent intended to stop the mugger as a means of getting his wallet back, but where a shotgun was aimed carefully at the mugger's back it is equally obvious that the *means by which* the agent intended to stop the mugger was to kill or at least gravely injure him. This fact cannot be avoided by some "inner act of 'directing [one's] intention,'" because to an important extent "circumstances, and the immediate facts about the means you are choosing to your ends, dictate what descriptions of your intention you must admit" (Anscombe, 2001, p. 63).

This suggests that Quinn needn't have been led by worries about *such* cases to shift away from traditional concerns about intending vs. foreseeing harm: we can apparently put a stop to absurd extensions of the idea of merely foreseen but unintended harm simply by attending more carefully to the structure of the cases, without even having to appeal to any controversial notion of closeness. In fact, however, this is too hasty: for there is one response open to the terror bomber that will in fact lead us right to the real problem of closeness that arises for other kinds of case as well, so that a complete answer even to this sort of case will require an account of excessive closeness. The problem is that the terror bomber might reply to the above argument by saying that his more proximate intended means of making the civilians appear dead is not *killing* them per se, but only impacting them sufficiently with the bombs to put them in a condition of appearing dead. He doesn't deny the intention of that proximate means, but only denies that their death or even their harm *as such* is intended as a means, though he admits it is foreseen as a side-effect of the intended means.⁷

This ought, I believe, to strike us once again as sophistry, but it is not as easily exposed and dismissed as the initial claim. A proper answer to this more subtle claim will require a careful look at the notion of excessive closeness, to flesh out and support the plausible idea that (i) people's being impacted with ordinary lethal bombs in such a way as to make them appear dead, and (ii) their being killed, are too closely related to allow conceptual room for the bomber to intend the first without intending the second (at least where he is aware of the relation

between the two, and not deluded about how the bombs work). Let us turn, then, to some familiar cases that highlight this issue more clearly, without the distraction of potential confusions between side-effects and proximate means.

3. EXCESSIVE CLOSENESS: THE BASIC ACCOUNT⁸

One classic example presents us with a large man hopelessly stuck in the exit of a cave, trapping his unfortunate fellow explorers inside: if they blow him to bits with some handy dynamite as a means to their end of clearing the exit and escaping, can they then claim that his death was a foreseen but unintended side-effect of their action – not only unintended as an end (since they presumably weren't aiming at his death for its own sake), but equally unintended as a means? Another familiar case involves a craniotomy abortion: if a doctor aims at crushing a fetus's skull as a means of removing the fetus in an emergency and saving the woman, can he then claim that its death was merely a foreseen but unintended side-effect – not only not intended as an end, but also not intended as a means?⁹ (See Foot, 1994a, p. 268 for both cases.) We may set aside the question whether so acting in these cases is ultimately morally justified for other reasons: our concern here is simply with the plausibility of these claims about intention. And as Foot has observed, to allow such claims would “make nonsense of [the intend harm/foresee harm distinction] from the beginning”, rendering it so flexible as to make it worthless for purposes of moral discrimination (Foot, 1994a, pp. 268–269). We thus need a principled way of explaining why its application cannot properly be extended in these ways.

It is not possible to dispense with these claims in the same way we originally sought to deal with the terror bomber's claim above (a move that proved to be incomplete in any event). In the cave case, for example, the agents grant that they intend to blow the man to bits as a means of clearing the exit, and we cannot dismiss their further claim not to intend his death by showing that they aim at his being killed as a *more proximate means* to bring about his being blown to bits: for it is not true that they

seek to blow him to bits *by* killing him (in the way that the terror bomber was originally accused of seeking to make people appear dead by killing them); the structure here is plainly different. Similarly with the craniotomy case: it is not as if the doctor accomplishes his means of crushing the fetus's skull *by* killing the fetus, so he is not guilty of mistaking more proximate means for unintended side-effects. Instead, the problem in these cases is the same problem we were left with after considering the terror bomber's potential response to the earlier objection – a problem concerning the splitting of hairs in ways that seem implausible, though it is not easy to explain why.

The problem, we might again like to say, is that the relation between what is admittedly aimed at as a means, on one hand, and the harm or death, on the other hand, is “too close” to allow for talk of aiming at the former without aiming at the latter (at least where the connection is known to the agent, as we shall assume throughout). For example, the relation between someone's being blown to bits, on one hand, and his being killed or his death, on the other, is simply “too close” to allow for talk of aiming at the former but not at the latter, as a means to one's end. The difficulty, however, is to give an account of what this means. Bennett has offered the most rigorous and interesting discussion to date of many of the natural possibilities. Yet while he does ultimately offer a very minimal and vague way of trying to answer the sophisticated terror bomber, he does not take this to provide much support for a robust use of the intend/foresee distinction, and he rejects all the more substantive proposals for understanding “excessive closeness” as clear failures (Bennett, 1995, pp. 203–213). There is, however, an important candidate that he has failed to consider properly, and it is this idea that I shall build upon. My proposal may be summarized as follows, to be developed and refined in Sections 4 and 5:

The relation between the intended means and the harm in question is “too close” to allow for application of the intend/foresee distinction when the relation between the relevant *states of affairs* is a *constitutive* one rather than a *merely causal* one.¹⁰

We may begin to elucidate this claim by first considering a simple case in which the relation is merely causal, and hence *not* “too close”, after which it will be easier to see what is different about cases in which the relation is instead constitutive. In turning a runaway trolley off of the crowded track on which it is traveling and onto a side track in order to minimize harm, an agent does not *intend* the subsequent death of the one person stuck on the side track, which is no part of her plan even if it is *foreseen*; the very presence of the one on the side track is entirely incidental, and his death is not in any sense being used as a means to accomplish the goal of getting the trolley turned away from the crowd and minimizing harm, though it is a foreseen side-effect of doing so (Foot, 1994a, p. 270).¹¹ Similarly, in the strategic bombing case mentioned earlier, the pilot targets a munitions plant, intending its destruction as a means to win the war, but does not similarly intend the foreseeable harm to nearby civilians from the fallout as a means to anything; that harm is no part of the agent’s plan, and the very presence of the civilians is incidental to it.

What creates the conceptual space for such claims about intention, so that they can be made without lapsing into sophistry? The key is that the relation between what the trolley driver aims at as a means, i.e. the trolley’s being diverted away from the crowd, and a person’s subsequently being killed down the side track, is *merely causal*: the trolley’s being diverted certainly causes the one down the side-track to be killed, but the trolley’s being diverted clearly isn’t *constitutive* of his being killed. We can thus clearly distinguish between the trolley’s being diverted and various *distinct results* of this, such as a person’s being killed down the side-track. And when combined with the other features of the case, this provides us with the conceptual space to say without sophistry that the agent aimed at the one thing (the trolley’s being diverted) but not at the other (the person’s being killed) as a means to her end. Similarly with a straightforward strategic bombing case (though we will also have to consider harder cases later): the factory’s being destroyed certainly *causes* the nearby civilians to be harmed from the fallout, but the factory’s being destroyed is plainly not

constitutive of their being harmed, and this is important to one's being able to claim that one aimed at a factory's being destroyed as a means *without* thereby aiming at neighboring civilians' being harmed as a means.

By contrast, a man's being blown to bits is not a *cause* of his being killed, as a Trolley's being turned onto a side track is a cause of the one person's being killed when it proceeds to roll over him. The detonation of the dynamite is the cause of the man's being blown to bits and of his being killed, but *his being blown up* is *constitutive* of *his being killed*. The people in the cave aim at his being blown up (as a means to clear the exit and escape), and the problem is that the constitutive relation between this and his being killed closes off any space for the intend/foresee distinction: they can't aim at his being blown to bits without aiming at his being killed, because this is not merely a matter of one thing resulting in another distinct thing (as the trolley's being turned resulted in the person's being killed), but a matter of the one thing's being constitutive of the other. One can no more aim at a man's being blown to bits without aiming at his being killed than one can aim (literally) at a spot on a target without aiming at the target it partly constitutes. Similarly with the crushing of a fetus's skull: the fetus's skull's being crushed *constitutes* its being seriously injured or killed, rather than just being a cause of it, so that it is impossible to aim at the first without thereby aiming at the second (except in the unlikely case of ignorance of the relation between skull crushing and injury or death, which we may set aside).¹² Again, we can contrast this with the trolley case, for even though someone is foreseeably crushed in the trolley case as well, the structure of that case is different: it is not the crushing we are aiming at as a means, as in the craniotomy case, but just the turning of the trolley away from the crowd; so there is not in that case the same problem of aiming at something (the crushing) that is too intimately related to the harm to avoid thereby aiming at the harm.

It is sometimes objected to the claim that the physician intends the death of the fetus in the craniotomy case that "it seems harsh to say that the physician desires the death of the

fetus”: surely she may “view the death as a tragedy” and desire it no more than does the physician who performs a hysterectomy on a pregnant woman in order to remove a cancerous uterus (Delaney, 2001, p. 570). Similarly, it will be said, in the cave case: the people in the cave do not desire the man’s death, as shown by the fact that they regret his death, and have no interest in seeking some more efficient way to kill him (such as lethal injection), but “confine their satisfaction [and efforts] only to the success of” the project of clearing the exit (Delaney, 2001, p. 575). This line of objection, however, misses the point, at least if meant as an objection to the claim that the deaths in such cases are intended rather than being merely foreseen but unintended side-effects. The fact that the physician regrets the death of the fetus in the craniotomy case no less than in the hysterectomy case, or that the people in the cave are not interested in other ways to kill the trapped man, shows only that the deaths are not desired *for their own sakes*, as *ends*. It does not show that they are not relevantly intended as means.¹³

The same point applies to counterfactual or hypothetical thoughts such as “if the cancer could be removed without bringing about the death of the fetus, then [the physician in both cases equally] would have absolutely no reason to bring it about” (Delaney, 2001, p. 571): this shows that the physician does not desire the death of the fetus as an *end*, but it leaves open the question whether she nonetheless intends it as a *means*. This is a point we have already seen in Section 2 about the limitations of counterfactual and hypothetical “tests”, in connection with the example of the terror bomber and his dispositions in certain miraculously altered situations. In the present, more complicated, set of cases, the point is that while the deaths are not aimed at as ends, there is something that (despite the counterfactuals and hypotheticals) is *in fact* clearly aimed at as a *means*, which is *too closely related* to the deaths for the latter to be regarded as mere side-effects in the course of pursuing the end. The physician, for example, obviously needn’t desire the death of the fetus for its own sake, but she does desire and intend the fetus’s skull’s being crushed. And the claim is that because this is *constitutive* of the fetus’s being

seriously injured or killed, rather than just being a cause of it, the physician cannot correctly categorize the fetus's being seriously injured or killed as a merely foreseen but unintended side-effect.

4. DEVELOPMENT AND CLARIFICATION OF THE ACCOUNT

There are several important points to make about this proposed account of the notion of closeness. The first is that it is important to attend carefully to linguistic constructions, and so to whether we are speaking of states of affairs or of acts. I have claimed that *the trolley's being diverted* and *the person's being killed* are distinct states of affairs, the former being a cause of the latter, rather than constitutive of it. But this is not to deny the plausible claim that *my diverting the trolley* is identical to *my killing the person on the side-track*, i.e. that there is *one act* here – perhaps a single movement of my hand – that falls under various descriptions, rather than a performance of many successive acts. (It's not as if I first turned the trolley and then later did something else that was the killing; rather, the various effects of my single action become incorporated as further descriptions of what I did (Anscombe, 1963; Davidson, 1980) – or so I shall assume here.) This thesis about act-identity is consistent with my altogether different claim about the distinctness of certain states of affairs, i.e. the trolley's being diverted and the victim's being killed, one of which is aimed at and one of which is not. Similarly, my claim about the constitutive relation in the cave and craniotomy cases is consistent with the Anscombe–Davidson claim but is not to be confused with it: *the doctor's crushing the fetus's skull* is *the doctor's killing the fetus* (the Anscombe–Davidson claim), but the distinct point I am making is that *the fetus's skull's being crushed* is constitutive (and not a mere cause) of *the fetus's being seriously injured or killed*.

My proposal thus avoids the difficulties Bennett has raised for the hopeless attempt to use the Anscombe–Davidson thesis of act-identity to explicate the idea of closeness. Such a move

would obviously be futile because it would wind up tying things too closely together even in cases where things are *not* plausibly “too close”. Since *my turning the trolley* just is *my killing the person on the side track*, for example, if we took this sort of identification to be what squeezes out any room for the intend/harm/foresee harm distinction, then we would disqualify the trolley case no less than the cave and craniotomy cases (see Bennett, 1981, pp. 108–109; 1995, pp. 207–208). Again, this is not a problem for my proposal, where the idea of being “too close” is understood in terms of a constitutive relation between states of affairs rather than in terms of the identification of the agent’s X-ing with the agent’s Y-ing. In some cases, such as the trolley example, the relation between the relevant states of affairs is merely causal, and there is room for the intend/harm/foresee distinction; in other cases, such as the cave example, the relation is a constitutive one, in which case there is not.

A second important point is that nothing we have said requires in general that the harm be at most likely but not inevitable if the intend/harm/foresee harm distinction is to apply.¹⁴ In the trolley case, for example, the driver might see with certainty that the one on the side track will be run over and killed, yet this does nothing to undermine her claim that it was merely a foreseen but unintended side-effect, entirely incidental to her plan: it simply underscores the “foreseen”. What matters is the relation between the intended means and the harm, and if this is merely causal and not constitutive, then the intend/harm/foresee distinction might apply even if the harm is foreseen with certainty. It is partly for this reason that I reject the suggestion (explored and also rejected by Bennett, 1995, p. 209) that *causal necessitation* is the relation that makes two states of affairs “too close” to allow room for aiming at one without aiming at the other: the inevitability of harm is not to the point. Similarly with the relation of *logical entailment*, which Bennett also considers and rejects: that would, among other things, capture too little, since a fetus’s skull’s being crushed does not logically entail its being seriously injured or killed, but is intuitively a central case of excessive closeness. This idea is properly captured by focusing on the fact that the

fetus's skull's being crushed is constitutive of its being seriously injured or killed.

There are admittedly some things we may *not* be able to say where the harm is foreseen as inevitable. It will not be possible, for example, to say that the harm to civilians from a bombing was an *accident* in a case where there was not merely a foreseeable risk of harm from a miss or from fallout, but virtual certainty of harm given the nature or strength of the bomb or the degree of proximity. And such actions will obviously be subject to criticism if insufficient measures were taken to protect civilians, which is what underlies criticism of the US Army's continued use of cluster munitions, for example (as in the latest war in Iraq), which pose special and ongoing dangers to civilians. Setting that problem to one side, however, the mere fact of inevitability of harm clearly does not turn the strategic bomber into the equivalent of the terror bomber, who targets civilians and aims at their harm as a means: whether the harm is literally "accidental" or not, it is at any rate not aimed at as a means or as an end.

This point, however, must not be pushed too far. Consider the case of *area bombing*, where a broad area containing both civilians and a targeted militant, for example, is destroyed as a means of killing the latter. Here again the harm to civilians is inevitable, but there is a difference from the strategic bombing case. The strategic bomber can honestly say that he was aiming (literally) only at the munitions plant and not at civilians, striving to hit it and not them, despite the fact that the fallout would inevitably kill those very nearby. In such a case, we still have a merely causal relation between the plant's being destroyed (which was his aim) and the resultant harm. By contrast, the area bomber cannot say that he is trying to hit only the militant and not the civilians: by using a blunt weapon that destroys the entire area, he leaves behind the possibility of any such discrimination, and this changes the structure of the case. His means to the end of killing the militant is to obliterate a large area containing him, which also happens to contain civilians, and so the question is how the obliteration of this area that contains both the militant and civilians is related to the

harm to the civilians. The answer, I believe, is that an area's being obliterated is *constitutive* of its occupants' being destroyed, making it sophistical to speak of aiming at the one without aiming at the other. That is, while the bomber's *proximate instrumental* means to kill the militant is to drop the bomb, this works *by* destroying an entire area, and an area's being destroyed is *constitutive* of its living occupants' being destroyed, much as a computer file's being erased is constitutive of the contained text's being erased.¹⁵

This serves as an answer to the "casuistical problem" that led Nagel (1979, pp. 59–63) at one point to shy away from the intend/foresee distinction, and explains why *such* cases of inevitable harm preclude its application. At the same time, to reject the application of the intend/foresee distinction in the case of area bombing, and to suggest that it is significantly more problematic than strategic bombing, is not to deny that area bombing is still morally different from terror bombing. One clear difference is that in area bombing – and likewise in "human shields" cases, where innocents are contained within a legitimate target such as a military bunker – the civilians' presence is entirely incidental: their welfare is being sacrificed in the pursuit of the end, but neither they nor their harm are being used as means, whereas in the terror bombing case their presence and harm are straightforwardly needed and used as means.¹⁶ The area bomber, like the strategic bomber and the one who attacks even in the face of human shields, acts *despite* the harm to civilians, while the terror bomber *seeks it out*. The latter behavior is likely to strike us, plausibly, as typically worse. Again, this is not to deny that the former may be seriously wrong, and all the more so if the act exhibits callous indifference to innocent life, as where the goal was of insufficient importance to justify such harm in any case.¹⁷ But where all else is held equal, it will clearly be at least as bad, and plausibly worse, deliberately to seek out the deaths of innocents with a strategic eye toward using them to further one's goals, treating innocent people "as if they were then and there *for* [one's] purposes" (Quinn, 1993b, p. 190).

I shall not here attempt to defend this intuitive moral claim, since such matters lie beyond the scope of the present investi-

gation. But it is worth noting the implausibility of one attempt to turn the tables here. Bennett (1995, p. 218) suggests that what the strategic bomber or area bomber (in contrast to the terror bomber) does “is in a way worse than treating [his victims] as means. He is treating them as nothing; they play no part in his plan; he is not *even* treating them as means.” This is surely unpersuasive. As Bennett himself notes, the agent needn’t be simply disregarding the value of his victims, “treating them as nothing,” but may take their value very seriously, acting *in spite of* it only after having carefully weighed it against what he takes to be the overwhelming importance of the military objective. It is hard to see how the fact that he is not further intending to exploit their deaths in the manner of the terror bomber can be used *against* him, as if somehow aiming at the deaths of innocents as a means to one’s ends shows greater respect for them.

A third point of clarification is that the proposed way of cashing out the issue of excessive closeness makes no appeal to the degree to which actions of the *kind* in question tend to be associated with the kind of harm in question. Quinn finds appeals to closeness uninviting because he assumes (following a suggestion from Hart) that it depends on such a distinction, the idea being that the relation between an intended means and a harm is too close if actions of that kind are “invariably connected with” harms of that kind – an idea he goes on to undermine with counterexamples (Quinn, 1993b, pp. 179–180). But this is a confusion. What allows the death in the trolley case to be a foreseen but unintended side-effect is not the fact that trolley turnings are not invariably associated with killings on side tracks, and what disqualifies the death of the fetus in the craniotomy case from being a merely foreseen but unintended side-effect is not the fact that fetal skull crushings are invariably associated with fetal deaths. What matters is just the nature of the relation between the intended means and the harm in the individual case – i.e. whether it is merely causal or constitutive. This focus avoids the difficulties Quinn raises.

Finally, it is important to note one limitation of the proposal. In all the examples I have considered, the relevant

relation between states of affairs – whether merely causal or constitutive – is simply a matter of *natural* fact. But there are also cases where the relation is determined by merely *conventional* arrangements involving the agency of others, and it may be much less clear what to say about such cases. For example, there may be a convention in place stipulating that a student who receives an “F” in any class is thereby ineligible for financial aid: having an “F” *constitutes* ineligibility for financial aid. Now I may or may not be aware of this policy. If I am unaware of it, obviously I can intend a student’s being given an “F” without thereby intending his being rendered ineligible for financial aid. But suppose I *am* aware of the policy (the same knowledge assumption we made in all the other cases). If we treat this case in the same way as the others, we must conclude that I cannot intend a student’s being given an “F” without thereby intending his being rendered ineligible for financial aid – for the relation is constitutive, not merely causal. But is that a plausible result? Even if I am aware of the policy, I may dislike it and regard it as unfair, in which case it might seem quite possible for me to intend the student’s receiving an “F” (because he failed my class and I am concerned to give him the grade he deserves) *without* thereby intending his being rendered ineligible for financial aid – this being a merely foreseen but unintended side-effect of my giving him the grade he deserves.¹⁸

It may be possible to handle this objection by just biting the bullet and defending the original implication. After all, the doctor who performs the craniotomy may well *regret* the death of the fetus, and *dislike* the connection between the procedure and the fetus’s death, but that does not undermine the conclusion that the death is intended in that case rather than being a mere side-effect. So it is not obvious that the grading case is really any different. Still, there is significant intuitive force in the thought that it is, so that it is worth considering whether we can give a principled explanation for why this case may be an exception, allowing for the intend/foresee distinction despite the constitutive relation between states of affairs.

A plausible answer is that the merely *conventional* status of the constitutive relation is what opens up the space for the intend/foresee distinction here, where there would not be such space in the case of a naturally constitutive relation; and it does so at least partly because of the involvement of others' agency. In the grading case, not only do I perhaps dislike the policy of making the presence of an "F" on a student's transcript constitutive of ineligibility for financial aid, but I can attribute its existence and enforcement to another agent, such as the university administration. There is a sense in which I can legitimately say: "I am simply giving you the grade you deserve for my class, and the rest is someone else's doing, which I would prevent if I could." By contrast, it would be absurd to try to say something parallel in the cave or craniotomy cases, claiming that I am simply blowing the man up or crushing the fetus's skull and the rest is someone else's (the universe's?) doing. There is no other agency in the natural relation cases, so that the agent is the only one to whom the harm is to be attributed. Thus, because of the constitutive relation between what the agent aims at and the harm, there is no room for the intend/foresee distinction. But in a merely conventional case, where the constitutive relation depends on conventions involving other people's agency, there is room for an agent to distance herself from certain potentially quite arbitrary upshots of the action. It is that distancing of one's agency from such upshots that creates conceptual space for the intend/foresee distinction here, despite the constitutive relation between the relevant states of affairs.

Let us, then, stipulate that the proposed account of "closeness" is to be restricted to cases in which the relation in question is *natural* rather than merely conventional. The claim, then, is that if the relation between two states of affairs is known to the agent, natural, and constitutive rather than merely causal, then we cannot properly speak of an agent's intending the one while merely foreseeing but not intending the other.

5. COMPLICATIONS AND HARDER CASES

The cases considered in Section 3 are relatively easy ones, at clear ends of the spectrum. In the cave case, for example, the agent plainly aims to do something foreseeably lethal *to the victim himself* as a means – namely, to blow him up – and his being blown up is *coincident with his death*. These factors make it natural to say not only that he is killed by being blown up, but that his being blown up plainly *constitutes* his being killed. Likewise with the craniotomy case. By contrast, in the trolley case, the victim's being killed is clearly distinct in time and place from the trolley's being turned, which latter doesn't itself involve the victim at all, but only an operation on an *object*. This makes it easy to say that the trolley's being turned is *not constitutive* of the victim's being killed, but only leads to it. But what if we modify the craniotomy or cave cases to move them closer to the trolley case in certain respects? Suppose, for example, that instead of the skull's being crushed, it need only be partially fractured or even just scratched, though this will foreseeably lead to complications eventually resulting in the death of the fetus. Or suppose what we need is not a man's being blown up, but only his being cut, though this will again foreseeably lead to his death. Someone might, say, cut an attacker with a knife in self-defense, where she aims at his being sufficiently injured to distract him long enough for her to get away, her only interest being in this temporary distraction, not in his death, though this may foreseeably follow later (perhaps because there is no medical aid available to stop the bleeding).

These cases are like the original craniotomy or cave cases in that some *injury to the victim* is plainly intended, unlike in the trolley case, where all that is intended is the turning of the trolley; but these new cases are unlike the cave and craniotomy cases and more like the trolley case in that the intended effect – the scratching or cutting – is *not coincident* with the death, and may even be quite removed from it. (Note that these cases are thus also very different from the terror bomber case, where the agent pursues his goal of making civilians appear dead *by* killing them. This is clearly not the structure of the present

cases, since the death of the attacker in the cutting-in-self-defense case, for example, doesn't even occur until *after* the relevant goal has been achieved.) What we have here is a structure similar to Bennett's example of the politician who deliberately destabilizes a union in December to prevent a strike at Christmas, foreseeing that this blow will lead ultimately to its permanent disintegration, but being interested only in the more immediate result of preventing the Christmas strike, not in the union's demise (Bennett, 1995, p. 210). What does the proposed account say about such cases?

When the man in the cave entrance explodes, it is relevant that *his dying* is not a separate *event* from *his exploding*, or at least is so nearly coincident to it as to make it entirely natural to say of the corresponding states of affairs that *his being blown up* is *constitutive* of *his being killed*, as we have done. By contrast, in the cutting-in-self-defense case, the man's dying – perhaps a full day later – is plainly a distinct event from the cutting, on virtually any plausible scheme for carving up events.¹⁹ We say, then, that *his being cut leads to his death* (though not that *his being cut leads to his being killed*, since that would imply that his being cut led to some other lethal event, as if it led him to walk in front of a car). And we do *not* say that his being cut *is* (or is constitutive of) *his being killed* (though we do say that *the agent's cutting him was her killing him*, this being the earlier point about *act-identity*). This is again because, unlike in the case of someone's being crushed or blown up, the *cutting* on Monday is plainly a distinct event from the *dying* on Tuesday, which makes the identification of *his being cut* with *his being killed* misleading, despite the *act-identity* noted above. "His being killed" refers here in a general way to the overall situation of which the cutting is a constituent, and of course we will say later that he was killed *by* being cut. But where the dying is so clearly a distinct event, as the man's dying is here clearly distinct from the cutting, it is not quite right to say that his being cut *is* or *constitutes* his *being killed*, but only that it *causes his death*. And this plausibly opens up conceptual space for the intend/foresee distinction: the agent can intend his being cut, as a means to her end, without thereby so intending his being

killed or his death. The sort of case in which she would intend *his being killed* as a means to her end would be a different one, where she deliberately shoots to kill, in order to immobilize him, in order to get away.

It will likely be objected here that while these cases are reasonably clear in the relevant respects, one can easily describe cases along the spectrum between the cave case and the cutting case, or between the craniotomy case and one in which a scratch leads to an infection and death days later. And this virtually guarantees gray areas where it won't be at all clear whether or not the dying is "a separate event" from the intended event, thus reintroducing obscurity into the account at the level of *events*. Indeed, it may seem in many cases to be more a matter of *degree* than a matter of the events simply being distinct or not, especially given the vagueness surrounding the very notion of the event of a person's dying. Isn't this, then, still a problem of the very same sort we began with? Have we made any real progress?

I believe that we have, and that the above considerations simply point to a limitation inherent in the nature of the subject rather than a deep problem for the account. We should fully acknowledge that there will be gray areas, as there are for most interesting topics. But just as the fact of dusk does not preclude our distinguishing day from night, the existence of gray areas in a philosophical account needn't undermine its interest or usefulness. It will do so only if the cases we want to have something to say about always lie squarely in the gray areas, so that the account proves incapable of providing guidance. This, however, is not the case with regard to the intend/foresee distinction, and we have made progress insofar as we can account for a range of philosophically interesting cases where the relevant factors are reasonably clear. The possibility of unclear cases does not detract from an account of the clear ones, such as straightforward cases of strategic vs. terror bombing. This is especially true if we understand why the unclear cases are unclear – what makes them pull in different directions. If a case lies along a spectrum somewhere between the cave case and the cutting case, we may be unable to place it neatly in one or the

other category. But if our account allows us to describe what matters about the case, what makes it fall where it does along the spectrum, and perhaps which end it falls closer to, then that is enough to be illuminating. The facts, and therefore any possible moral implications, may be no more determinate than that, and an account should not aspire to outdo the facts with artificially imposed precision.

The upshot for the above cases is that the proposed account allows – correctly, I believe – that one can in principle intend such injuries as a cutting, scratching, partial fracture or local disruption *without* thereby intending the death or demise that eventually results, even if the latter is foreseen. If the agent's interest is only in a certain local effect – such as distracting the attacker just long enough to get away, or temporarily disrupting union operations to prevent a Christmas strike – then it is plausible that clearly distinct further events may be merely foreseen but unintended, and the account preserves this result. This does not, of course, guarantee that the action will be morally justified. Whether the bringing about of foreseen but unintended harms is justified depends on many further factors. In the case of self-defense, it depends on whether the cutting was reasonably believed to be necessary to ward off the attack, and whether the foreseen harm was proportional to the harm being defended against. Similarly, we can recognize that the politician did not intend the demise of the union as a means or end, while still condemning his action on the grounds that he had no business doing something that would foreseeably destroy the union. As Foot has argued, initiating a harmful sequence of events will often be wrong, even if the harm is only foreseen but not intended, since this will often violate negative rights and duties (Foot, 1994a, b). We can thus allow the application of the intend/foresee distinction in cases with this structure without fear that this will automatically justify too much.²⁰

In this section, we have so far been discussing hard cases that share one feature of the earlier clear cases of excessive closeness – namely, that the agent aims at doing (or allowing) something foreseeably lethal *to (or for) the victim himself* as a means – while departing from the earlier cases insofar as the death of the

victim is *no longer roughly coincident with* the intended event. I have argued that this departure is often enough to open up sufficient conceptual space for the intend/foresee distinction, which is at least sometimes plausibly of moral significance and so should not be papered over (as it is by Quinn's shift in focus, discussed in Section 6). Let us now turn to the reverse kind of hard case: what if the death of the victim *is* roughly coincident with the plainly intended act, as in the cave and craniotomy cases, but we depart from such cases insofar as the foreseeably lethal involvement of the victim is now *irrelevant* to the agent's purpose, making it in this respect more like the trolley case?

This is the structure of the familiar hysterectomy case, where a physician removes a pregnant woman's uterus as a means of preventing the spread of cancer, foreseeably resulting in the death of the fetus. The involvement of the fetus is not itself intended as a means, as the fetus's presence is entirely incidental to the physician's purpose: she would have proceeded exactly as before, removing the uterus, had the fetus not been present. This makes the present case clearly different from the craniotomy case even though the death may be similarly coincident with the plainly intended event (the removal of the uterus). I believe it is again plausible to treat this variation as allowing conceptual space for the intend/foresee distinction, so that the death of the fetus is not intended as a means, though it is a foreseeable consequence of the intended removal of the uterus (as a means to contain the cancer).²¹ The relation is not "too close" here because the fetus's dying is plainly a distinct event from the uterus's removal, despite being roughly coincident in time with it, and this in turn blocks the claim of constitution for the relevant states of affairs. Whereas in the other cases, such as cutting-in-self-defense, the relevant distinctness of events was made clear by temporal differences (someone's being cut on Monday is not the same event as his dying on Tuesday), it is here brought out by the difference in the *objects*: the removal of the uterus is plainly not the same event as the dying of the fetus.

As before, then, while it is true that removing the uterus of a pregnant woman *causes* the *death* of the fetus, we should resist any move to the claim that the uterus's being removed

constitutes the fetus's being killed in the way that its skull's being crushed does. What it *constitutes* is the initiation of a sequence leading quickly to the death of the fetus, but that is a looser relation, and again allows conceptual space to speak of intending the removal of the uterus as a means to the end of stopping the spread of the cancer, without aiming at the death of the fetus as a means. Compare: I need a plank for something, but someone is standing on it, so that when I pull it out he falls and is injured. I intend the plank's being removed so that I can use it elsewhere, and the plank's being removed leads immediately to his injury, but *the plank's being removed* is not *constitutive* of *his being injured* (as his being blown up would be). According to the account, then, it is possible for me to aim at the plank's being removed as a means to my end *without* thereby aiming at his being injured as a means at all, his presence being incidental to my purposes. There will probably be other considerations (such as my likely violation of his negative rights) preventing the consideration about intention from doing much *justificatory* work in this case, but the point is that the intend/foresee distinction can nonetheless apply in cases with this structure, and may at least *sometimes* be morally relevant.

6. CAN WE SIDESTEP THE ISSUE OF CLOSENESS?

I have not attempted to provide a reductionist theory giving necessary and sufficient conditions for a relation's being constitutive rather than merely causal, because I do not believe any such theory to be possible without imposing unhelpfully artificial precision on a subject matter that is inherently messy. As Quinn points out in another context, "almost no familiar distinction that applies to real objects is clear in all cases, and theoretical reducibility is a virtue only where things really are reducible" (Quinn, 1993a, p. 157). Indeed, it is noteworthy that Bennett himself ultimately downplays the significance of the existence of fuzzy borderline cases, emphasizing instead his claim to have uncovered a much deeper difficulty that is "less like twilight than like a blazing sun in a black, star-studded

sky” (Bennett, 1981, pp. 113–114). What he means here is his alleged exposure of the “absurd results” to which our use of the concept of *intending something as a means* gives rise. I have tried to show, however, that he has not succeeded in this: there is no “blazing sun” here.

Still, some may prefer to avoid reliance on this less than fully analyzed distinction between merely causal and constitutive relations, at least if there is a simpler alternative that can equally rescue moral uses of the intend/foresee distinction from arbitrariness and absurdity. It is thus worth examining more closely one prominent suggestion along these lines, to see whether it really provides a satisfactory alternative.

In an effort to avoid the complications surrounding the appeal to closeness, Quinn (1993b, pp. 183–188) has proposed shifting the focus away from harm itself, so that what is discriminated against is not just intending harm as a means or end, but intending someone’s *involvement* in some harmful event, as a means or end (i.e., “direct agency”). It has been pointed out that this is ambiguous between (1) intending someone’s involvement in some event, as a means or end, where such involvement will in fact be harmful to the person (whether this is known to the agent or not), and (2) intending someone’s involvement in some event, as a means or end, where it is known to the agent that such involvement will be harmful to the person (Fischer et al., 2001, p. 193 f.). I believe it is clear from the kind of rationale Quinn gives for his view, however, that it is the second, much more plausible, interpretation that is intended. Let us, then, restrict our focus to that.

The advantage of Quinn’s shift in the focus of the intend/foresee distinction is that we no longer need to try to determine whether “the harm itself” is intended – i.e. as being “too close” to the intended means – or merely foreseen: it is enough for “direct agency” that a certain involvement is intended as a means or end and that this involvement is foreseen to be harmful to the victim (except in cases where it is independently within the rights of the agent to bring about such involvement, as discussed below). This is a significant move, and it is worth

pointing out that at least some prominent arguments against Quinn's proposal fail.

Bennett, for example, offers as a counterexample a case in which we quarantine a group of patients with a highly infectious disease while we await the arrival of the special masks that will allow the doctors safely to treat them – though we foresee that after a week the disease will almost certainly have progressed too far for them to be cured. Now intuitively, we do not intend their deaths, and this case thus belongs with the easier-to-justify cases of merely foreseen harm, such as the trolley case. But Bennett thinks that Quinn will be forced to treat it as a case of “direct agency,” because we “intend to “involve” the infected people” – even, Bennett thinks, if we “quarantine” them simply by removing ourselves (Bennett, 1995, pp. 212–213). This is mistaken, however, for the simple reason that on Quinn's view we have direct agency only where the intended involvement is what leads to the harm for those involved, and while it is true that we intend the patients to be isolated from us, it is plainly not *that* “*involvement*” that kills them: what kills them is their disease, and we certainly did not intend their continued involvement with the disease as a means to some end of ours. The latter would be the case in Quinn's example of “Guinea Pig”, where we deliberately allow a disease to progress in a patient for research purposes, but the present case is obviously not like that.

Another of Bennett's alleged counterexamples fails for equally clear reasons. He imagines a case of heroic self sacrifice, where a man ties a line to himself and swims out toward a ship where the line is needed to rescue the sailors, though he foresees that given the storm and the rocks he will almost certainly perish in the process (Bennett, 1995, p. 220). The worry is that Quinn will be saddled with condemning such an action as involving “direct opportunistic agency”. But such a worry is misplaced because this is not a case of direct opportunistic agency at all. We have direct agency, on Quinn's view, only where the involvement in question involves a violation of the independent moral rights of the person involved, and this is not the case where an agent voluntarily

embarks on an heroic mission. Of course, if we changed the case so that the man was forcibly tied to the line and set out into the currents, *then* we would have a case of direct opportunistic agency. But this is exactly what we should want to say in that case.²²

There are, however, some good reasons for doubting whether Quinn's proposal is satisfactory. He will, for example, have to categorize cases such as area bombing as examples of "indirect agency", since the involvement of civilians may be irrelevant to the bomber's purpose, the civilians' very presence being altogether incidental, such that "we act exactly as we would if they were not there" (Quinn, 1993b, p. 187). Quinn will thus have to group this case with cases such as trolley as far as intention is concerned. Yet this seems wrong, for reasons given earlier. Similarly, Quinn's response to a clever example from David Lewis is plainly inadequate, as Kamm has pointed out, and this produces an unattractive result (Kamm, 1992, pp. 377–378).²³ We can see this more clearly by modifying the example slightly.

Suppose someone wishes to demoralize the enemy government during a war, and knows he could do so by creating the belief that large numbers of civilians have been killed. He is reluctant to bring this about by deliberately killing civilians, but thinks he sees a way around this. It turns out that a major population center has been evacuated *unknownst* to the central government, so that if it is bombed they will *believe* that there have been massive casualties, simply viewing it from afar. He thus plans to go ahead and bomb the empty city. At the last minute, however, he discovers that the population has returned. Now suppose he goes ahead with the mission anyway. Quinn would like to say that the bomber "strictly intends to involve [the civilians] in something ... in order to further his purpose precisely by way of their being involved," thus likening the case to ordinary terror bombing, as seems plausible. But Quinn cannot really make this out, because although the bomber foresaw the harmful involvement of civilians in this case, *their involvement was nothing to his purpose*: all that is needed in this case is the government's seeing the destruction of the city from afar, for which the population's presence and hence involvement was entirely irrelevant (which is why the bomber was prepared to

destroy the city when it was empty). Thus, Quinn's account does not give the plausible answer he wants, and it seems that the only way for him to get it would be once again to fall back on some notion of *closeness* – just as I have proposed. For it is precisely the idea of excessive closeness here that disallows regarding the deaths of the civilians as merely foreseen side-effects of the destruction of the city, for reasons similar to those in the area bombing case, and so groups this case properly with ordinary terror bombing cases, despite its special features.²⁴

Finally, Quinn's account has the disadvantage of failing to explain – in a way that employs the intend/foresee distinction – the cases from the previous section. His version of the DDE discriminates against intentional involvement of someone as a means where this involvement foreseeably leads to his death, even if the death itself is unintended. Thus, in order to justify the cutting-in-self-defense case, he will have to appeal solely to *other* considerations, such as the attacker's having forfeited his usual right not to be killed; the intend/foresee distinction will not do any justificatory work here, since we are going *against* the thrust of its justificatory force on his view. By contrast, the account I have offered can explain why the act is justified at least *in part* by appealing to the intend/foresee distinction itself, stressing the fact that the death is not intended. Even if the other considerations would be sufficient to justify the act, the fact that the death is not intended as a means seems to be relevant, making it at least *easier* to justify than cases where the death is plainly intended as a means – where, for example, one deliberately brings about the death of an attacker as a way of immobilizing him and thus achieving the goal of getting away. Quinn's account misses this.

For the above reasons, then, I believe we cannot avoid dealing head on with the problem of closeness if we wish to preserve both non-arbitrary and intuitively plausible appeals to the intend/foresee distinction in normative ethics. The account I have developed and defended, in terms of the distinction between *merely causal* and *constitutive* relations among states of affairs, is intended to provide just such an account of excessive closeness, answering the challenges raised by Bennett and others in a principled and intuitively plausible way.

NOTES

¹ Anscombe (2001) refers to this general claim as the “principle of side-effects” to distinguish it from specific versions of the DDE that spell out the circumstances under which the bringing about of foreseen but unintended harm is permissible. I shall be concerned only with the general claim in speaking of the DDE.

² This is a point overlooked by Foot (1994a). She there argued that one could apparently do all of the work typically assigned to the doctrine of double effect with (roughly) the doing/allowing distinction, cashed out in terms of the distinction between negative and positive duties. She later broadened her view to give a place to both distinctions in normative ethics.

³ Bennett was once quite dismissive here, giving up almost immediately on the idea “in the absence of any help” with its meaning (1981, pp. 107–108). His more recent treatment of “the tight binding problem” (1995, p. 204 f.), however, is much more searching and extensive, though I shall argue that it is still unsatisfactory. Others are more sympathetic but likewise troubled enough by difficulties surrounding intention to marginalize the intend/foresee distinction in ethics. Nagel, for example, was once led by just such problems to shift his focus to a different distinction around which he thought a deontological constraint could be formulated with fewer difficulties (1979), as described in Section 4 below. He later, however, returned to the intend/foresee distinction as a central factor in normative ethics (1986, p. 179 f.).

⁴ There are, of course, other important problems that also need to be addressed. Thomson (1999, pp. 509–518) and Rachels (1994), for example, press an objection to the DDE that would be devastating if it worked. I have argued elsewhere, however, that their criticism is based on a mistaken construal of the DDE (FitzPatrick, 2003b).

⁵ The view he is attacking is the “entailment proposal”, according to which the relation between the plainly intended means and some effect is “too close” if and only if the latter is logically entailed by the former (1995, pp. 209–210) (see Section 4 below).

⁶ Kamm (1996, p. 155) briefly mentions a similar response to Bennett, which is developed along lines similar to what follows in the text by Delaney (2001, pp. 577–578). In Section 6, I will examine a more complicated and interesting variation on this example.

⁷ I thank Anubav Vasudevan for emphasizing this possible reply.

⁸ This section and the next build upon an idea that is sketched, though not developed or defended in detail, in FitzPatrick, (2003a). Parts of these sections borrow closely from that preliminary discussion to set up the more developed treatment here.

⁹ H.L.A. Hart once made such a claim, denying that the craniotomy case is any different in terms of intention from the case of an emergency hys-

terectomy performed on a pregnant woman. Bennett (1981, pp. 105–109) says something similar. I discuss both cases below.

¹⁰ Delaney (2001, pp. 568–572) considers and dismisses a view that somewhat resembles this proposal, though he conflates it with a different proposal involving an undeveloped notion of essential relations (the *relata* of which are also specified differently in different places), and does not address the important issues brought to light by Bennett (1995) in formulating it – issues that make a crucial difference, as discussed in Section 4 and later sections. Some of his objections are answered below; others are effectively answered in the course of developing my account with Bennett’s more developed challenges in mind.

¹¹ We may here ignore the fact that Foot does not believe that we need to rely on the intend/foresee distinction in justifying the decision to turn the trolley. Thomson (1986) likewise offers an account that does not require appealing to the distinction between intending and merely foreseeing harm. I am not here making any general claims about when that distinction is or is not crucial to moral justification.

¹² It is, unsurprisingly, not hard to modify relatively easy cases like this to make them much less clear, eventually shifting them into the other category. See Section 5 below.

¹³ It would, of course, be incorrect to say that “the physician intends to bring about the death of the fetus by crushing its skull,” as Delaney (2001, p. 571) points out. But again, that is just because such a construction would imply that the death is intended as an *end*, which it is not.

¹⁴ In this I depart from Anscombe (2001, p. 64), and also from Bennett (1995, pp. 224–225), who does ultimately concede a *small* potential moral role for the intend/foresee distinction, but only where the harm is merely likely rather than inevitable.

¹⁵ An alternative account for at least some cases would be that the area bomber seeks to kill a certain person (who might not be individually identifiable) by killing everyone in a containing group (e.g. a village). In that case it is even more clear that the deaths of the innocents are intended as a means, the death of the group being the means to the death of a given member.

¹⁶ The human shields case is a complicated one, falling somewhere between area bombing and strategic bombing. Attacking a military bunker that also contains civilians is like area bombing and unlike strategic bombing insofar as one *cannot* claim here to be aiming (literally) at the military target but not at the civilians; but it is unlike area bombing and more like strategic bombing in that one *can* claim to be targeting just the military target as such, rather than a wider area that includes it along with civilians. This may well make a moral difference, especially when combined with the partial shifting of responsibility to those responsible for putting the civilians in harm’s way as human shields. I shall set these interesting issues to one side here.

¹⁷ This is the sort of criticism that was raised, for example, in connection with Israel's bombing of an apartment building in Gaza City in July, 2002, killing 14 civilians along with targeted Hamas militant Sheik Salah Shehada.

¹⁸ I am grateful to Eugene Mills for raising this objection, and to other members of the audience at the Virginia Philosophical Association Meetings, 2003 for helpful discussion of these issues.

¹⁹ I shall not here attempt to provide or to defend a general account of the individuation of events. For my purposes it is enough that the claims about events are consistent with plausible existing accounts, though someone could of course raise further objections based on alternative accounts of events. Note also that I bring in events not as the primary category for solving the problem of excessive closeness, but only secondarily as described in the text, to explain why some cases plainly do *not* involve a constitutive relation among the relevant states of affairs. My account is therefore not directly vulnerable to the problems Bennett raises for an account of excessive closeness appealing simply and generally to event identity (1995, pp. 205–207).

²⁰ It is worth pointing out that where it *does* justify in cases of this type it will run counter to the spirit of Quinn's revisionist account. This is discussed in Section 6 below.

²¹ Again, I do not mean to imply that this claim is necessary to justify the act: the operation might well be justified independently, by appeal to other factors such as the lesser moral status of the fetus. Our present concern is simply with the application of the intend/foresee distinction.

²² For more on the rights-condition in Quinn's view, see Fischer et al. (2001, fn. 8).

²³ Oddly, Kamm does not mention that Quinn discusses this sort of case in a footnote and attempts to deal with it (see below), but her discussion does show why his answer fails.

²⁴ Fischer et al. (2001, p. 196, 201) likewise claim that Quinn will be forced to fall back on some notion of closeness, though for different reasons that I find less compelling.

REFERENCES

- Anscombe, G.E.M. (1963): *Intention*, 2nd ed., Oxford: Blackwell.
- Anscombe, G.E.M. (2001): 'Medalist's Address: Action, Intention and 'Double Effect'', in P.A. Woodward (ed.), *The Doctrine of Double Effect*, Notre Dame: University of Notre Dame Press, pp. 50–66.
- Bennett, J. (1981): *Morality and Consequences*, The Tanner Lectures on Human Values II, Salt Lake City: University of Utah Press, pp. 45–116.
- Bennett, J. (1995): *The Act Itself*, Oxford: Oxford University Press.
- Davidson, D. (1980): *Essays on Actions and Events*, Oxford: Oxford University Press.

- Delaney, N. (2001): 'To Double Business Bound: Reflections on the Doctrine of Double Effect', *American Catholic Philosophical Quarterly* 75(4), 561–584.
- Fischer, J.M., Ravizza, M. and Copp, D. (2001): 'Quinn on Double Effect: The Problem of 'Closeness'', in P.A. Woodward (ed.), *The Doctrine of Double Effect*, Notre Dame: University of Notre Dame Press, pp. 189–210.
- Foot, P. (1994a): 'The Problem of Abortion and the Doctrine of the Double Effect', in B. Steinbock and A. Norcross (eds.), *Killing and Letting Die*, New York: Fordham University Press, pp. 266–279.
- Foot, P. (1994b): 'Killing and Letting Die', in B. Steinbock and A. Norcross (eds.), *Killing and Letting Die*, 2nd ed., New York: Fordham University Press, pp. 280–289.
- FitzPatrick, W.J. (2003a): 'Surplus Embryos, Non-Reproductive Cloning, and the Intend/Foresee Distinction', *The Hastings Center Report* 33(3), 29–36.
- FitzPatrick, W.J. (2003b): 'Acts, Intentions, and Moral Permissibility: In Defense of the Doctrine of Double Effect', *Analysis* 63(4), 317–321.
- Kamm, F. (1992): 'Non-consequentialism, the Person as an End-in-Itself, and the Significance of Status', *Philosophy and Public Affairs* 21(4), 354–389.
- Kamm, F. (1996): *Morality, Mortality Vol. II: Rights, Duties and Status*, Oxford: Oxford University Press.
- Nagel, T. (1979): 'War and Massacre', in *Mortal Questions*, Cambridge: Cambridge University Press, 53–74.
- Nagel, T. (1986): *The View From Nowhere*, New York: Oxford University Press.
- Quinn, W. (1993a): 'Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing', in *Morality and Action*, Cambridge: Cambridge University Press, pp. 149–174.
- Quinn, W. (1993b): 'Actions, Intentions and Consequences: The Doctrine of Double Effect', reprinted in *Morality and Action*, Cambridge: Cambridge University Press, pp. 175–193.
- Rachels, J. (1994): 'More Impertinent Distinctions and a Defense of Active Euthanasia', reprinted in B. Steinbock and A. Norcross (eds.), *Killing and Letting Die*, 2nd ed., New York: Fordham University Press, pp. 139–154.
- Thomson, J.J. (1986): 'The Trolley Problem', in *Rights, Restitution and Risk: Essays in Moral Theory*, Cambridge, Harvard University Press.
- Thomson, J.J. (1999): 'Physician-Assisted Suicide: Two Moral Arguments', *Ethics* 109, 497–518.

Department of Philosophy
Virginia Tech
Blacksburg, VA 24061-0126
USA
E-mail: wfitzpat@vt.edu