

*Scanlon's Investigation: The Relevance of Intent to Permissibility**

Surely, one might think, intent matters morally. If I hurt you, the morality of what I did depends on what I meant to do. Was it an accident, or was it deliberate? If it was deliberate, did I have a good reason (was I trying, say, to remove a splinter)? Or did I do it out of spite? In *Moral Dimensions*, T.M. Scanlon agrees that intent may affect *some* moral dimensions of what one does, the idea that there are several moral dimensions being perhaps the theme of the book as a whole.¹ For example, it can affect whether it is, or is not, appropriate for you to blame me for what I did. But in the first three chapters of the book, he questions whether intent affects the *permissibility* of what one does, at least in the ways that many have naturally thought. There is, of course, much else to discuss in this rich book, not least the account of blame in the fourth and final chapter. But, in part because that account has attracted searching scrutiny elsewhere,² and in part because of limits of space, this review essay focuses just on his investigation into the relevance of intent to permissibility.

1. Interpreting the Conjecture:

It's easy to think that Scanlon's investigation of the relation of intent to permissibility aims to vindicate the following:

* I am grateful for comments on drafts of this paper from Frances Kamm, Véronique Munoz-Dardé, Jay Wallace, and, especially, Tim Scanlon. Many of the ideas were developed in my and Wallace's graduate seminar on *Moral Dimensions* at the University of California, Berkeley in the spring of 2008. I owe many debts to participants in that seminar, including, in particular, Sam Scheffler and Gerald Dworkin. Some sections of the paper were presented at an Author Meets Critics session at the Eastern Division Meeting of the American Philosophical Association in December of 2009. I benefited greatly from Scanlon's response and from the discussion that followed.

¹ T.M. Scanlon, *Moral Dimensions: Permissibility, Meaning, Blame* (Cambridge: Harvard University Press, 2008).

² See, in particular, R. Jay Wallace, "Dispassionate Opprobrium," and Susan Wolf, "Blame, Italian Style," in Samuel Freeman, Rahul Kumar, and R. Jay Wallace, ed., *Reasons and Recognition: Essays on the Philosophy of T.M. Scanlon* (New York: Oxford University Press, forthcoming).

Conjecture: When the fact that one acts for a certain reason helps to explain why one's action is impermissible (or permissible),³ it does so only derivatively.

Although Scanlon himself doesn't offer as definite an interpretation of "derivatively," I take it that the fact that one acts for a certain reason helps to explain why one's action is impermissible *derivatively* just when that fact helps to explain it only in virtue of constituting a morally relevant consideration, or satisfying the condition of a morally relevant principle, that could have been constituted, or satisfied, by a different kind of fact: that is, a fact other than that one acted for a certain reason. Scanlon freely accepts, for example, that it might be the case that calling one's aunt out of genuine concern for her welfare would be permissible, whereas doing it only to keep oneself in her will would be impermissible. This is because the latter action involves deliberately misleading her about something: namely, that one is genuinely concerned for her welfare (p. 40). Here the morally relevant principle might be:

If one knows that P is not true, it is impermissible to lead someone to believe P.

³ One might add here: "in such a way that one's action would not have been impermissible (or permissible) if one had done it but not for that reason." Scanlon tends to pass over the question whether the reason for which an action is performed merely strengthens the case for its impermissibility (or permissibility), without changing the ultimate verdict. Instead, he focuses on the question whether it strengthens the case to such an extent that the action, which would otherwise be permissible, becomes impermissible (or vice-versa). See, for example, p. 13: "the difference between causing harm intentionally and doing so negligently, however, is not a difference in *permissibility*. Both are generally impermissible." This leaves untouched the question whether malicious intent makes the case for the impermissibility of the first *stronger* than the case for the impermissibility of the second, without affecting the verdict. See also p. 39: Principle M characterizes a certain kind of wrong. An action would not be wrong *in that way* if it did not involve intentionally misleading someone. But it would not follow that what the agent would then be doing was permissible. The other features of action described in Principle M might still make it wrong. So, although the idea of acting intentionally plays a role in characterizing a particular form of morally objectionable action, it does not in this case play a role in drawing the line between what is permissible and what is impermissible.

In this case, P is that one is calling out of concern for her welfare, and the fact that one knows that P is not true is a fact about one's reasons. But in another case, say on a tax return, P might be that one had no investment income, and knowing *that* was false would have nothing to do with one's reasons. Similarly, it might be the case that refusing to renew a tenant's lease simply in order to free the apartment for one's daughter would be permissible, whereas doing so chiefly in order to get rid of the tenant would be impermissible. This might be the case if one had (with odd self-absorption) promised never to refuse a renewal except in order to free, from pure parental love, the apartment for one's daughter (p. 65). Here the relevant principle might be:

If one has promised to refrain from X-ing unless some condition obtains, then it is impermissible to X unless that condition obtains.

In this case, the condition is that one is not renewing the lease *in order to* free up the apartment for one's daughter, and so the fact that the condition obtains is a fact about one's reasons. In another case, however, the condition might be that one's friend got another lift to the party, which might release one from a promise to give him a lift (in one's already cramped car) if he couldn't manage any other way of getting there. The fact that this condition obtains would not be a fact about one's reasons.

This Conjecture is natural to attribute to Scanlon in part because he deals with the vast majority of the relevant cases by arguing that, despite first appearances, they conform to it. Later, I will suggest that he holds something weaker, as is most clearly indicated by his handling of certain cases in ways that do not conform to the Conjecture. At any rate, there may be some risk of distortion in representing Scanlon as specially concerned to vindicate *any* general view of the relation between intent and permissibility, rather than simply investigating whatever relations there are and letting the chips fall where they may. But it does help to organize discussion of

Scanlon's investigation, and to bring out some of the more important points, to see how far this Conjecture can be pushed.

2. Meaning and the critical use of principles:

One finds several different kinds of arguments in *Moral Dimensions* that might be taken to support the Conjecture. First, Scanlon distinguishes the moral dimension of permissibility from a second moral dimension, which he calls "meaning," of which blameworthiness is one species. This can be used to block an objection to the Conjecture: namely, that since actions done for bad reasons are morally objectionable in some way, they must be impermissible. What may be objectionable, Scanlon suggests, is the meaning of such actions, not their wrongness.

Consider what we might call "harmless malice" cases, where someone does something otherwise permissible, but for a bad reason. A man sticks pins in a voodoo doll in order to make his wife suffer, when, as he should know, it won't have this effect. There seems something morally objectionable about this, and this seems to be precisely his reason for doing it. (There would be nothing morally objectionable, although perhaps something pointless, about his sticking pins in the doll in order to *relieve* his wife's suffering.) To this, Scanlon can reply that our sense that there is something morally objectionable about the action reflects not a judgment about its permissibility, but instead a judgment about its meaning (p. 46).

How does permissibility differ from meaning? We consider questions of permissibility characteristically when we are trying to decide what to do, or in helping someone else to decide what to do. Thus, questions of permissibility find their home, one might say, in the deliberations of the *agent prior* to reaching a decision. By contrast, we consider questions of meaning characteristically when we are trying to come to terms with what someone's decision implies for his relations with others. Thus, questions of meaning find their home in the assessments that

others make *after* the decision has been made.⁴ In assessing the meaning of someone’s decision to X, we are not *ourselves* deciding whether to X, or indeed whether to do anything else—although the conclusions drawn from an assessment of meaning may itself bear on further decisions, such as whether the wife is to live with her husband as before. In these assessments of meaning, the reasons for which the agent acted are of the first importance. When the question is what the wife is to make of her relationship with her husband, it matters crucially whether his voodoo was meant to harm or succor her. By contrast, when the question is his deliberative question, “May I cast this spell?” it is less clear that it matters what his own reasons are, as opposed to what the likely effects of casting the spell will be.

The contrast between permissibility and meaning is part of a broader contrast between “deliberative” questions, which have their home in first-person deliberations of the agent prior to decision, and “critical” questions, which have their home in third-person evaluations of the agent after the decision. Deliberative questions include not simply what I may permissibly do, but also what I ought to do or believe more generally. Critical questions include not simply what the meaning of someone’s action is for his relations to others, but also whether the action is morally worthy in a more impersonal sense, or skillfully executed, or reasonably taken. One of the important, if comparatively understated, contributions of Scanlon’s book, I think, is to clarify and call attention to this basic distinction. It is significant, I believe, for a host of other

⁴ Speaking of “the deliberations of the *agent prior* to reaching a decision” and “the assessments that *others* make *after* the decision has been made” helps, I think, to indicate the relevant distinction. But it may be a ladder that one kicks away after ascending. The more one thinks about the distinction, the less confident one is of any neat correspondence with these standpoints. After all, we can make claims about the reasons that third parties have, purely hypothetically or after the fact. And we can consider before the fact, indeed in deliberation, how we ourselves would properly be evaluated for a particular performance. Compare p. 218 n. 9. I am indebted to Sam Scheffler for pressing this point.

philosophical questions, including the nature of rationality and the relativity of “ought” claims to evidence.⁵

3. Choosing to Act for a Reason:

In addition to “harmless malice” cases, there are “helpful malice” cases, in which an agent does something otherwise *required*, but at least in part for a bad reason. In Scanlon’s example, the agent “hates the person who needs help and would be happy to see him die, but she saves him anyway because she does not want him to die *right then*, since that would mean that his heir, with whom she is locked in a bitter political contest, would have much more money to spend on his campaign” (57). Again, the distinction between meaning and permissibility allows Scanlon to agree that there is some moral criticism of the agent’s conduct, without having to claim that it is impermissible. But he also presents a more direct argument—this is the second sort of argument—that the agent does not act impermissibly.

First, it is implausible that it is permissible not to save the person at all. Second:

Always a Permissible Option: There must be at least one permissible option in any given situation.

And, third, that there are only two candidate permissible options: not saving and saving for the bad reason. What about a third candidate: saving for the good reason? With regard to the more unusual case, in which the agent does not believe that there is any such reason, Scanlon rules this out by appeal to:

May Implies Can Choose: If one cannot choose to Y in a given situation, then Y is neither a permissible, nor an impermissible option for one in that situation.

⁵ See Niko Kolodny, “Aims as Reasons” in Freeman, Kumar, and Wallace, ed. *Reasons and Recognition*; and Niko Kolodny and John MacFarlane, “Ought: Between Subjective and Objective,” unpublished.

and:

Believing is Necessary: If one does not believe that R is (sufficient) reason to X, one cannot choose to X for R.

Since the potential savior doesn't believe that there is the good reason, he cannot choose to save for the good reason, and so it can't be permissible.

With regard to the more usual case, where the agent *also* believes that there is a good reason, Scanlon rules out the third candidate by appeal to *May Implies Can Choose* and

Believing is Sufficient: If one believes that one has reason, R, to X, then one cannot choose to X but not for R.

Since the potential savior believes that there is the bad reason, if he chooses to save, he necessarily chooses to save for the bad reason. So he cannot choose to save for the good reason alone, and so it can't be permissible for him to save for the good reason alone.⁶

Why accept *Believing is Necessary* and *Believing is Sufficient*? Perhaps these principles rest on a view about what it is to do something for a reason. Begin with the:

⁶ *May Implies Can Choose* and *Believing is Necessary* and *Sufficient* supply a quick argument for the claim that:

there is no situation in which one's X-ing for reason R is permissible (or impermissible), but in which one's X-ing but not for R is impermissible (or permissible).

Suppose one permissibly X's for R (e.g., calls one's aunt out of concern for her welfare). Then by *Believing is Necessary*, one believes in that situation that R is a reason for X. Then by *Believing is Sufficient*, one cannot choose in that situation to X but not for R. Then by *May Implies Can Choose*, it is not impermissible to X but not for R in that situation.

However, this claim is not the Conjecture. The Conjecture is a claim about whether intent *explains* the impermissibility of an action in a situation, not a claim about whether *alternative actions* with different intents *are* permissible in that situation. A more appropriate test of the Conjecture is not whether, in the actual situation, it *was* permissible, say, to call one's aunt for out of genuine concern for her welfare, but instead whether, in the "closest" counterfactual situation in which one *did* call for a different reason, doing so *would have been* permissible. And if one had X-ed but not for that reason, then presumably one would have chosen to do so, with whatever beliefs were necessary for so choosing. So *May Implies Can Choose* does not rule out the possibility that one would have done so permissibly.

Minimal View: One X's for R just when one believes that R is reason to X and one X's. Then, if one does not believe that R is a reason to X, then one cannot X for R. And if one cannot X for R, then one cannot choose to X for R. So we get Believing is Necessary. The Minimal View implies that if one believes that R is reason to X and one X's, then one X's for R. So one cannot X but not for R. So one cannot choose to X but not for R. So the Minimal View implies Believing is Sufficient.

However, suppose we accept instead the:

Conditional Intention View: Choosing (to X for R) just is forming a conditional intention (to X insofar as R supports X-ing),

and understand X-ing for R as X-ing from such an intention. This is a way of putting a point made by Alec Walen.⁷ It seems plausible that:

one can intend (to X insofar as R supports X-ing) only if one believes that R is a reason to X.

And this gives us Believing is Necessary. But it is unclear why we should accept Believing is Sufficient. One might believe that R supports X-ing, but choose to X but not for R: by, for example, not intending (to X insofar as R supports X-ing), while intending (to X insofar as R' supports X-ing).⁸ This pattern of intention would be substantially different from choosing to X for R. Suppose that, before executing the intention, one came to believe that a certain factor had arisen that undermined the support of R' for X, but not the support of R for X. Then, unless one formed some other intention, one would not X. By contrast, if one had intended (to X insofar as R supports X-ing), then one would X (insofar as one was rational) without forming some further

⁷ Alec Walen, "Choosing Your Reasons for an Action" (unpublished). However, Walen doesn't speak of conditional intentions.

⁸ This would not be irrational, if one believed that R was merely a sufficient reason (to X for R).

intention. If Believing is Sufficient is not true, then it isn't clear why, in the more usual case, in which agent is not so wicked as not to recognize the good reason to save the person, the agent can't be required to save the person for that reason alone.

So long as Believing is Necessary, May Implies Can Choose, and Always a Permissible Option remain, Scanlon can still argue at least that, in the more unusual case, in which the agent does not recognize the good reason, he is not required to save the person for it. This does raise a problem, however, if we also accept:

Must Choose for a Reason: If one chooses to X, then one chooses to X for some R.

Then we get the result that:

For any action X, X-ing is permissible in a situation in which one does not believe that there are sufficient reasons to refrain from X-ing, and refraining from X-ing is permissible in a situation in which one does not believe that there are sufficient reasons to X.

For example, take the unusual case in which the agent does not see the good reason for saving the life, but alter it so that she no longer sees even the bad reason for saving the life. Believing is Necessary implies that there is no R such that the agent can choose to save for R. Then Must Choose for a Reason implies that the agent cannot choose to save. By May Implies Can Choose, it is not permissible for the agent to save. By Always a Permissible Option, there is a permissible option that does not involve saving. So the agent is permitted not to save. It is one thing to say that any intuitively impermissible action becomes permissible for anyone who is *not in a position to know* that there are sufficient reasons not to do it. But it is quite another to say,

as this consequence does, that it becomes permissible for anyone who simply does not believe it, even if he *is* in a position to know better.⁹

Perhaps, though, Scanlon could reject Must Choose for a Reason. Sometimes, he might say, people choose irrationally, for no reason at all.¹⁰ In this unusual situation, therefore, morality can forbid the agent from choosing to act on the reasons as he sees them and instead require him to choose for no reason at all. However, if this route is taken, then in the case in which the agent sees only the bad reason for saving the person, while the agent cannot be required to save *for the good reason*, the agent can be required to save *but not for the bad reason*: that is, required to save for no reason at all, something which the agent can choose to do. (Of course, one might wonder what substantive moral argument can be given for this requirement. The point is that Scanlon's reflections on what agents can choose would not rule it out.) Suppose this is so. Then in any saving case, either the agent believes that there are only bad reasons for saving, in which case he is required to save, irrationally, for no reason at all; or the agent believes that there is also some good reason for saving, in which case he is required to save for that good reason alone (since Believing is Sufficient has been abandoned); or the agent does not believe that there is a bad reason for saving, in which case it cannot be permissible to save for a bad reason (since Believing is Necessary remains). In other words, it turns out that it is *never permissible, no matter what the agent believes, to save for a bad reason*. So Scanlon's argument seems not to establish its intended conclusion.

4. Cases: Predictive significance

The bulk of Scanlon's inquiry into the significance of intent for permissibility—and the third kind of argument for the Conjecture—is more substantively moral. Scanlon examines a variety

⁹ I'm indebted here to conversation with Jay Wallace.

¹⁰ In comments, Scanlon raises this possibility.

of cases in which the reason for which one acts appears to make the action nonderivatively permissible or impermissible, and he offers alternative explanations of these appearances.

Scanlon's explanations are deceptively rich and subtle. Although I have some questions about his analyses, much of my aim in what follows is just to spell out, clearly and compactly, what I take them to be.

To begin with, Scanlon observes that a bombing sortie may be permissible if done *merely* in order to achieve a legitimate military objective, whereas the sortie would be impermissible if done *also* to terrorize the population. This is because the latter intent will make the bomber prone to conduct the sortie in ways that will produce unjustifiable harm to civilians. Here, the bomber's intent makes the sortie impermissible. But it does so only derivatively (pp. 12–13, 30).

The relevant principle is that:

It is impermissible to X if the conditions are such as to make it likely that one's X-ing will cause unjustifiable harm.

In this case, the relevant condition is that the bomber intends to terrorize the population. But in another case, the condition might be that the medicine that one is about to give the patient is contraindicated.

In this case, one's intent has what Scanlon calls "predictive significance": it is evidence about what one's action is likely to be or bring about. And evidence about what one's action is likely to be or bring about, Scanlon holds, can affect its permissibility. This is because, for Scanlon, permissibility is determined neither by what is in fact the case, nor by what the agent believes is the case, but instead by what the evidence available to someone in the agent's position suggests is the case.

I believe that the idea of predictive significance is powerful, with implications beyond those that Scanlon draws here.¹¹ And I agree with Scanlon's case against the "objective" ought, which is supposed to depend only on the facts, and which has been championed by G.E. Moore and Judith Jarvis Thomson. The basic point is that deliberative notions such as "ought" and "permissible" need to be able to guide agents in their decision-making; that is, so to speak, their function or purpose. But agents typically need to make decisions without knowing all of the relevant facts. So if judgments about what they ought or may do require such knowledge, then such judgments cannot guide their decision-making.

However, I think that it is a mistake to say, as Scanlon does, that deliberative notions such as 'ought' and 'permissible' depend on the evidence *available to someone in the agent's situation*. This is to embrace a version of the "subjective" ought (even if less "subjective" than an ought that depends on what the agent believes or, if this can come apart from the evidence available to someone in the agent's situation, the evidence he actually possesses). Thomson's arguments against the subjective 'ought' are *also* compelling. Suppose we know that B's flipping a switch will cause great harm to A, although the evidence available to someone in B's situation presently suggests that it is harmless. "Wouldn't it be weird in us," Thomson writes, "knowing what will happen if B flips the switch, to say, 'Look B, we know something you don't

¹¹ First, I suspect that the phenomena that seem to suggest putative, autonomous rational requirements on intentions—that they be consistent and means-end coherent—can largely be explained by appeal to the predictive significance of intention. When one intends something, this makes it likely that one behave in certain ways, and this fact about one's likely future, just like facts about what others will do or what will happen regardless of what anyone will do, can give one reasons to intend or not to intend certain other things now. See my, "The Myth of Practical Coherence," *European Journal of Philosophy* 16 (2008): 366–402, and "Instrumental Transmission," unpublished. Second, I believe that predictive significance helps to explain how intending an end can give us reason for the means, most notably in the "tie-breaking" cases described by Scanlon, "Reasons: A Puzzling Duality?" in R. Jay Wallace, Philip Pettit, Samuel Scheffler, and Michael Smith, ed., *Reasons and Value: Themes from the Moral Philosophy of Joseph Raz* (Oxford: Clarendon Press, 2004), pp. 231–246.

know. If we tell you, then it will be true to say that you ought not flip the switch, but not if we don't tell you" (48). Scanlon's response is puzzling. The remark, he writes, is "odd because it seems to suggest that it would be B's knowing about the harm, rather than the harm itself, that would make it the case that he ought not flip the light switch." But one would have thought that this was Thomson's point. The view that 'ought' is subjective implies, weirdly, that if B does not know (or, as might suit Scanlon's account better, if B is not *in a position* to know) about the harm, then it is not the case that B ought not to flip the switch, whereas if B does know about the harm (or is in a position to know about the harm), then B ought not flip the switch. In this sense, the view that 'ought' is subjective implies that B's knowing (or being in a position to know) about the harm, rather than the harm itself, makes it the case that he ought not flip the switch. Scanlon continues: "If B knows about the harm, then he should count *it*"—rather than his knowing about it—"as a reason against turning on the light" (51). But how can B count the harm, whether or not he knows (or is in a position to know) about it, as his reason, when the subjective view of 'ought' compels him to accept that he ought not flip the switch *if and only if he knows* (or if and only if he is in a position to know) about it?

In any event, the spirit of Thomson's objection can be expressed in another way, to which Scanlon's response seems not to apply. Suppose we know we can't make our evidence available to B (and couldn't make it available to anyone in his situation). We can't communicate with him, or he has reason to mistrust us. Wouldn't it be *not* "weird in us," wouldn't it be natural, to say to ourselves, or just to think: "B ought not flip that switch! It's likely to do great harm to A!"? If our 'ought' were subjective, however, then we would know that what we say is false. Relative to the evidence available to someone in B's situation, it's not true that B ought not flip the switch, despite the fact that, as we know, it's likely to do great harm to A.

Scanlon accepts that there is *also* an “objective” use of ‘ought,’ which depends on what would be best when all the facts are in, and which lacks the “the moral content”—or, one imagines that he would also say, the deliberative role—“that the idea of permissibility has” (p. 48). So he might argue that this is the ‘ought’ that our advice, or would-be advice, uses. But I don’t think this can be right either. It would mean that we were talking past A, or changing the subject, rather than advising him: that is, helping him arrive at the correct answer to the questions that he is deliberating about: “What ought I to do? What may I permissibly do?” What we say seems to have as much moral content and deliberative relevance as those questions themselves. So if one finds it odd to say, “B’s not morally permitted to flip the switch, since it will cause A great pain, which gives him grounds for objection that outweigh any reasons that B has to do it,” I feel that one is tacitly hearing this remark, *pace* Scanlon, as implying something about B’s blameworthiness, or some other critical, rather than deliberative, dimension.

The dilemma, then, is this. On the one hand, ‘ought’ and ‘permissible,’ as well as ‘what is likely’ cannot depend on all the facts, since they need to be able to guide the agent’s decisions. On the other hand, they cannot depend solely on the agent’s epistemic situation either, since that closes off advice from better-informed advisors. When we ask, “What ought I to do? What may I permissibly do?” we take these questions to be ones on which other people, with different epistemic positions, might appropriately weigh in. Somehow, therefore, these deliberative notions have to be neither objective nor subjective, but rather something in between. Of course, it is a further question how this can be.¹²

5. Cases: The appeal of the Doctrine of Double-Effect

¹² See Kolodny and MacFarlane, “Ought: Between Subjective and Objective” for further criticism and a positive proposal.

Scanlon's interest in the relevance of intent to permissibility derives in large part from his doubts about the doctrine of double effect, which says, roughly, that it is permissible to bring about some bad when one *merely foresees* that this is necessary to produce a sufficiently greater good, but impermissible when one *intends it as means* to that good. For example, Tactical Bombing—inflicting a certain number of civilian casualties as the foreseeable effect of destroying a munitions plant—seems permissible, whereas Terror Bombing—inflicting the same number of civilian casualties in order to demoralize the population—seems impermissible. Likewise, Drug Shortage—using the available medication to save five people with the foreseeable effect that a sixth person, deprived of it, dies—is permissible, whereas Drug/Transplant—withholding the available medication from one person so that her organs will become available to save the five—is impermissible, even though both cases involve saving five while letting one die.

Is there any way to explain these distinctions without assuming that intent nonderivatively affects permissibility? One approach would be to assign moral significance to, as it were, “external” correlates of intending and merely foreseeing. What matters is the *abstract type of relation*, causal or otherwise, between the lesser evil and the greater good that the action does (or, better, can be expected to) bring about, whether or not the agent *intends* these relations. For example, it might be said that in Tactical Bombing, the lesser evil (i.e., the death of civilians) is only an *effect* of the greater good (i.e., the destruction of the munitions plant), whereas in Terror Bombing, the lesser evil *causes* the greater good (i.e., the weakening of public support).

First, one might doubt that such relations can bear the moral weight assigned to them. Indeed, I suspect that Scanlon doubts this; elsewhere he questions whether the fact that an action would use one person's death as a *means* to save five is sufficiently important to give him a decisive objection to the action, an objection which he would lack to an action that made his

death merely a *consequence* of saving the five (pp. 118–121). Why should the one care sufficiently about this difference in the relation of his death to the saving of the five?

Second, it is hard to see how any such account of the relations would be able to sort the relevant cases appropriately. For example, suppose that in Tactical Bombing the destructive force that eventually destroys the munitions plant must first to pass through the space where the civilians are, thereby killing them. It would seem no less permissible than the case in which the explosion of the munitions plant causes the civilian deaths. And yet the same abstract causal relation would seem to obtain here as in a case in which, rushing to the hospital to save five, we drive through the space where the one is standing, thereby killing him, which is intuitively impermissible.¹³

Scanlon's approach is different and, in a way, more elementary. He simply notes that in each case there is an underlying principle that forbids certain actions except under certain conditions, which do not depend on intent. While these conditions may obtain in Tactical Bombing and Drug Shortage, they do not obtain, or it would be absurd to suppose that they do, in Terror Bombing and Drug/Transplant.

Applying this schema, he suggests that while there may an exception to the impermissibility of inflicting civilian casualties, this exception requires at least that the casualties “can be expected to bring some military advantage, such as destroying enemy combatants or war-making materials” (28). Whereas the destruction of a munitions plant is plausibly a military

¹³ For an attempt, pursued with singular ingenuity and intellectual honesty, to spell out morally relevant abstract relations of this kind, see Frances Kamm, *Intricate Ethics* (New York: Oxford, 2007), esp. Ch. 5. At pp. 173–176, she discusses the difficulty of explaining the permissibility of natural versions of Tactical Bombing by appeal to the relations that she identifies, and she makes something like point just made in the text. Kamm suggests that the permissibility of Tactical Bombing might be explained instead by more permissive principles that apply to harming innocent members of an enemy nation. This may be so, but it would not, in itself, explain the distinction between Tactical and Terror Bombing: why the latter is *impermissible*.

advantage, “the death of noncombatants is not rendered a ‘military advantage’ by the fact that it would shorten the war by undermining public morale” (29). Presumably, a third case, “Mixed Tactical Bombing,” in which bombing can be foreseen to destroy the munitions plant, but one authorizes the bombing, with moral opportunism, in order to cause civilian casualties, would also be permissible on this account (assuming that the concerns about predictive significance discussed in the previous section could be set aside).

I doubt, however, that the notion of a “military advantage,” or any other restricted kind of *advantage or objective*, will sort the cases in the right way. To begin with, it doesn’t help to say that *the death of noncombatants* is not a “military advantage.” What we need, in order for Scanlon’s sketched principle to deliver the right result, is that the *analogue to the destruction of munitions*, namely, *weakening public support for the war effort*, which the death of noncombatants “can be expected to bring” about, is not a military advantage. But one might wonder why not, especially when one considers the military’s seemingly legitimate, if ham-handed, involvement in PsyOps and “winning hearts and minds.” Moreover, wouldn’t it be permissible—no less permissible than Tactical Bombing—to drop packets of very persuasive leaflets where this, unavoidably, would cause civilian casualties on impact? Even if undermining *public* support is not a military advantage, undermining *combatants’* will to fight seems like a military advantage if anything is. So how will the sketched principle handle a Terror Bombing that aims (as presumably the firebombing of Dresden and Tokyo on some level aimed) to get soldiers to lay down their arms, or their commanders to accept terms of surrender? More to the point, what if the aim of Terror Bombing is to convince workers at the munitions plant to rise up and to destroy it? By hypothesis, the destruction of munitions certainly counts as a legitimate military advantage. The moral, it seems, is that the relevant principle must

discriminate not only in the intended objective or foreseen advantage, but also in the way in which civilian deaths are causally related to it. If so, can the appeal to the doctrine of double effect, or to an “externalized” counterpart, really be avoided?

Nevertheless, I still feel that Scanlon’s underlying logic, in its broad outlines, does explain the fundamental difference in our reactions between Tactical and Terror Bombing. That logic, as I understand it, takes as its starting point a dialectical position in which it is assumed that there is some significant prohibition on killing noncombatants. Whatever its precise contours, the prohibition looks something like this. We may cause combatant casualties, provided this meets standards of proportionality; the advantage for our side must be sufficiently valuable to justify their losses. But we may not, save for highly restricted exceptional conditions, cause noncombatant casualties. Doing so is simply off the table, prior to, and preempting, any consideration of “proportionality.” It is not simply, for example, that a greater advantage must be achieved to justify noncombatant casualties than to justify combatant casualties. The question is then what these exceptional conditions are. Crucially, can an opportunity to give enemy combatants or noncombatants an *incentive*, or to otherwise *psychologically pressure* them, to weaken their war effort or to aid ours qualify as an exceptional condition, which then allows us to enter noncombatant casualties into the calculus of proportionality?

It is hard to see how the answer could be yes. The problem is that such opportunities are pervasive. Killing civilians *always* provides such an incentive, *always*, by amplifying the horrors of war, applies such psychological pressure—perhaps not to a Moscardó or Zhukov, but no doubt to the less denatured soldiers under their command, not to mention the civilians behind their lines. To count an “opportunity” for civilian deaths to work some effect on the minds of the

enemy as an “exceptional condition” would be in effect to make civilians fair game, subject *as a rule* to the calculus of proportionality. And this is precisely what the present dialectic already assumes is not the case.

To be sure, the difference between Terror and Tactical Bombing, on this account, turns on how civilian deaths are causally related to a given advantage or objective. But what is problematic about the causal relation in Terror Bombing, on this account, is not that civilian deaths are intended as a means, nor is it that civilian deaths are causally “upstream” from the relevant advantage. It is instead that the civilian deaths bring about the advantage by way of an effect that civilian deaths *always* have, and so that a principle that allowed us to enter civilian deaths into the calculus of proportionality *whenever* they had that effect would be indiscriminate. There’s no appeal to intent here, or to some intrinsic disvalue in certain kinds of causal relations. There’s only an appeal to the assumption that there is a distinction between combatants and noncombatants of the familiar kind, which, by default, prohibits noncombatant casualties, prior to any consideration of proportionality.

Turning to the drug cases, we have a principle of the form:

One is required to see to it that condition C (e.g., that the one is alive) holds (by, e.g., giving life-saving aid to the one), unless one could achieve G (e.g., giving life-saving aid to five others) by not seeing to it.

This exception is what makes it permissible to distribute the medication to the five in Drug Shortage. However, there is another principle of the form:

If condition C holds (e.g., the one is alive), one is required not to X (e.g., to harvest his organs) even if one could achieve G by X-ing.¹⁴

This is why one is not permitted to harvest the one's organs while he is still alive in Transplant. Scanlon notes that the justification for withholding the medication in Drug/Transplant would have to take the following, "absurd" form:

The agent is permitted not to see to it that C holds, because if C did *not* hold, the agent would be *permitted* to X and thereby to achieve G.

I share the sense that this justification is defective in some way. But one is left wanting some more explicit articulation of the defect, which the book does not give.

In personal communication, however, Scanlon suggests to me a way to bring out the absurdity. If such justifications were acceptable, then the more extensive a person's claims that we not treat him in certain ways (e.g., harvest his organs), the more moral reason we would have not to meet claims that he would otherwise have (e.g. to be kept alive). A person's claims on us would thus be at odds with one another. The more reason we have to recognize claim *A* despite the good that we could do by ignoring it, the less reason we have to recognize would-be claim *B*, since by failing to recognize *B*, we might be liberated of the need to recognize *A*.

6. Cases: Using as a means

Kantian moral theorists have paid great attention to actions that seem to be impermissible because they:

(a) involve someone as a means to one's end,

¹⁴ See p. 33 for this exception. This isn't quite the principle that explains Scanlon's Transplant, but it might as well be. It would explain a more brutal version of Transplant, in which one skips the lethal injection and just proceeds to harvest organs from the live person.

which is a matter of the intent with which they are performed. Now, Scanlon thinks that not all such actions are impermissible. For example, mutually beneficial, consensual cooperation is not. Do such actions become impermissible when they also:

(b) “subvert the person’s will,” by depriving her of a say about her involvement?

Not quite, Scanlon suggests, since some actions with features (a) and (b) are permissible. In the *Paris Case*, A plans to go to Paris. A doesn’t want B to go to Paris. A knows that if B knows A’s plans, then B will go to Paris. So A does not tell B. If it is *permissible* for A to refrain from telling B that A is going to Paris, Scanlon says, then A need not have done anything impermissible. However, *if it was impermissible* for A to refrain from informing B that A was going to Paris, Scanlon says, then A *has* done something impermissible.

What has A impermissibly done? An obvious answer is that A has impermissibly failed to inform B. Although Scanlon does not articulate it, the principle that would support this answer would seem to be something like:

Preventing Informed Decision about Involvement: If A can inform B that B’s action (e.g., not going to Paris) risks involvement in something that B has reason not to want to be involved in (e.g., missing the chance to be in Paris with A, strained though it is to call this a kind of “involvement”), then this gives B some claim on A to inform him of this, which, when weighed against the competing claims of A, may make it impermissible for A to fail to inform B of this.

However, it is also possible that A has also wronged B in a further way by *going to Paris without B*. The relevant principle would then be:

Involving Having Failed to Inform: If A has impermissibly failed to inform B in the way described by Preventing Informed Decision, then it is impermissible for A to bring about the relevant involvement (e.g., by going to Paris, given that B will not go).

The reason why the Paris case intuitively counts as one in which A “uses” B, or treats B “as a means,” is that *A seeks to benefit from B’s involvement* (i.e., his absence from Paris). But A’s seeking this does not *make* what A does *wrong*. It would be wrong in the same way if A gave no thought at all to B’s involvement.

There are, however, two kinds of case in which the fact that B’s involvement benefits A does make what A does wrong: that is, where this fact helps to fulfill the relevant conditions for impermissibility. As an example of the first sort of case, A wants B to attend his talk because B’s presence will impress A’s prospective employer. A knows that B has reason to avoid attending the talk if it will have this effect, so fails to inform B. A, we might imagine, wrongs B by failing to inform B. (Perhaps A also wrongs B by giving the talk, having failed to inform B.) In this case, the fact that B would be *furthering A’s end* is the involvement that B has reason to want to avoid, and so triggers the requirement to inform B. Hence, it *makes* the action impermissible. But in another case the undesired involvement might be furthering *someone else’s end*, or something independent of furthering anyone’s ends. So the fact that B’s involvement benefits A makes the action impermissible only *derivatively*.

In the second sort of case, A wants B’s involvement for instrumental reasons, which give it a particular kind of meaning: making it a case in which A is “just using” B. For example, A wants B to go with him to the dance just to appear to be in the “in crowd” and not for romantic reasons. B has reason to want to avoid interaction with this meaning. So it is impermissible for A to fail to inform B that his involvement would have this meaning: that is, to fail to inform B of

A's reasons for wanting B's involvement. This again seems to be explained by Preventing Informed Decision (and perhaps also Involving Having Failed to Inform), where the involvement that B has reason to want to avoid is "just being used by" B.

The careful reader will have noticed that I have not appealed to the principle that Scanlon officially invokes to explain these cases:

Costly Involvement Without Consent

A X's impermissibly if:

1. "A's [X-ing] affects or involves B in a way that is impermissible unless B consents to it.
2. Under the circumstances, B lacks the opportunity to give or withhold consent," whether or not A has brought this about (p. 110).

Such a violation occurs, for example, when one farmer diverts jointly owned water from another's fields without consulting him.

I have no objection to this principle as such. And indeed it gives clearer guidance than the alternative principles that I have proposed, since it does not require any further balancing of the claims of A and B. The trouble is that I don't see how it applies to the Paris or job-talk cases. For Scanlon, recall, A does something impermissible in the Paris case just when A is required to inform B that A is going to Paris. At times Scanlon seems to suggest that:

if A is required to inform B that A will, by X-ing, involve B in something that B has reason not to want, then it is impermissible for A to X unless B consents.

For example: "the question is whether... [this] makes it the case that his consent is required. Based on the analysis that I have offered, this depends on the claim that the person has to be informed" (119). If this were true, then the Costly Involvement principle would apply to the

Paris case. But it seems false. Suppose A has promised to let B know if A ever goes to Paris, so that B will have the chance to go with A (i.e., to avoid involvement in A's going without B), although A has made it clear that A has not promised to B not to go without B. A is required to inform B, but so long as A does inform B, A is free to go without B's consent. Suppose A *does* inform B that he is going to Paris next month, but B cannot go then and, in any event, B tells A in no uncertain terms that he does not want A to go without him. B has not "consented" to A's going without him. But A is still permitted to go to Paris without B.

Likewise, in the job-talk case, is there some action of A's that is impermissible without B's consent: e.g., A's giving the talk in B's presence when this will further A's job prospects? Of course, B is free to prevent A from performing this action by not showing up. But is this the same as saying that A needs B's consent in order to perform it? Suppose B, for whatever reason, can't leave the room. Does he then have veto power over A's giving the talk? Not in the most natural elaborations of the case. At any rate, would the case be interestingly different if there were *no* relevant "involving" action on A's part (other than the action of failing to inform B), and so no involving action of A's requiring B's consent? Suppose A inadvertently catches sight of B about to sign up, under the watchful eyes of A's prospective employer, to attend A's talk the following day, and A fails to inform B of the implications, although he could easily do so. One might say that B is "unable to consent to" an action of *B's own* (namely, B's involving himself: helping A to get the job, by being seen to sign up for the talk) in the sense that he is unable to make an informed decision to do it. But this isn't the sort of "consent" relevant to Costly Involvement, since it is not consent to an action of A's.¹⁵

¹⁵ Moreover, in the Job Case, Scanlon writes that: "The wrong in question... is that A wrongfully deprives B of information" (113), which suggests that the *only* wrong is that of failing to inform. If the "the" is meant literally, then there is no further wrong of *involving* B (such as diverting the

These complications arise, I think, because there are two ways in which A may be responsible for B's involvement in something. First, A can do things that lead B to do something that involves B. Such is the failure to inform in the Paris and job cases. Here it doesn't make sense to speak of B's consenting to A's involving B, since what is at issue is B's involving himself. Second, A can do something that involves B. Such is the involvement without consent in the water diversion case, or the involving having failed to inform in the Paris and job cases. Here it may make sense to speak of B's consenting to A's involving B, although it may be that B has an objection to what A does even if B's consent is not required.

In any event, Scanlon's main point stands. Whether A counts as "using" B "as a means" depends on whether A seeks B's involvement as a way to further A's ends. But this fact plays no (nonderivative) role in explaining why what A does is impermissible under any of these principles.

7. Cases: Incentive Threats

Threats, along with warnings, are announcements in which:

X tells Y that if Y does A, C will happen, believing that Y does not want C to happen.

Threats are distinguished by the fact that:

- (i) C is an intentional action of X's,
- (ii) X announces an intention to C if Y does A, *because* X believes Y does not want C to happen,

Incentive threats are distinguished by the further fact that:

- (iii) X makes the threat in order to get Y not to do A.

water from B's fields in the farmer case), which is what the Involvement Principle is supposed to explain. (Note that this also suggests that Involving Having Failed to Inform does not apply. Giving the talk, after having failed to inform, is not itself impermissible.)

Thus, incentive threats challenge the Conjecture.

One response would be to deny that actual intent to influence plays any role in explaining why the announcement is impermissible. Take an intuitively impermissible incentive threat. Would it be permissible to do the same thing if one did so *not* in order to influence the hearer? (Of course, such an announcement would not count as an “incentive threat,” but that is a different question.) Suppose someone makes the same announcement as a lame joke, when he should realize that anyone who hears it will interpret it as a genuine incentive threat. Intuitively, the announcement is still impermissible. And the principles that Scanlon identifies to explain the impermissibility of incentive threats seem to bear this intuition out. Actual intent to influence the hearer plays no role in fulfilling any of the conditions of the relevant principles.

In any event, Scanlon takes a different approach. First, he notes that X may wrong Y by making Y’s choice situation worse by causing him to *expect* the *penalty*, C. X wrongs Y in this way just when it is impermissible of X to C. What makes it wrong is that:

it is impermissible to give Y grounds to fear that X will do something, when X’s doing it is impermissible.

The conditions of the relevant principle can be satisfied by announcements do not seek to influence the hearer. Non-incentive threats and warnings can wrong Y in precisely the same way. So if intent to influence makes the announcement impermissible (and, as we have seen, it may be only the appearance of such intent that matters), it does so at most derivatively.

Second, X may wrong Y by making Y’s choice situation worse by changing the *meaning* of Y’s response: e.g., making it cowardly, etc. Presumably, it isn’t always wrong to worsen other’s choice situation in this way, although Scanlon does not give any general rule for telling when it is or is not wrong. In this case, it seems less plausible that it is wrong *only* when it is

impermissible to C. Offhand, it seems that I might wrong you by telling you that unless you stop smoking, I won't read your drafts, even though I'm under no obligation to read them.

Non-incentive threats and warnings can wrong Y by changing the meaning of Y's response. Admittedly, incentive threats—and “offers,” in which C is something that Y otherwise has reason to want—can change the meaning of Y's response in *distinctive* ways, precisely because of (the appearance of?) the intent to influence. Such (appearance of?) intent can make Y's response servile, a case of being pushed around, and so on. Nevertheless, intent still plays at most a derivative role, if we think of the underlying principle as roughly that:

It is impermissible for X to change the meaning of Y's response *in ways in which Y has reason not to want* (if these changes are sufficiently important as weighed against the relevant claims of X).

Scanlon then goes on to consider a special class of incentive threats in which:

(iv) Y's not doing A=Y's not interacting with X in a way that is permissible *only if Y consents*.

The paradigm is: “Your money or your life.” Scanlon observes, first, that Y's engaging in that interaction (i.e., handing over the cash), under threat of C, *when C is impermissible*, does not amount to consent, and that such threats seem permissible just when C is permissible. Now Scanlon doesn't say *why* such threats (as opposed to the subsequent actions: either imposing C, or engaging in the interaction without genuine consent) are impermissible just when C is impermissible. Perhaps it comes under the principle suggested earlier that it is wrong to lead someone to fear that one will do something wrong: in this case, the fear would be, so to speak, of a disjunction of wrongs: that one will either impose C, or engage in the interaction without genuine consent.

In any event, Scanlon argues that, since such threats are impermissible just when C is impermissible, the impermissibility of such threats depends on X's reasons only if the impermissibility of C depends on X's reasons. He acknowledges cases in which the permissibility of C *does* depend on the agent's reasons. For example, it seems wrong to fire an employee *for* refusing sexual advances, when it would otherwise be permissible to fire the employee. Again, this is not, in the first instance, a way in which an *announcement* can be wrong, but instead a way in which a certain *penalty*, a certain C, can be wrong. It would be wrong to fire the employee for refusing sexual advances, even without any prior announcement.

Scanlon suggests, plausibly, that this is wrong because it involves the employer having “an unacceptable form of control over” the employee. But can this be squared with the Conjecture? Can we identify an underlying principle that makes no reference to the agent's reasons? The relevant principle seems to be that it is wrong to withhold goods from people *in certain ways* when the power to withhold goods in those ways would constitute an *unacceptable form* of control. The question is whether “in certain ways” can be understood in some way broader than simply “for certain reasons,” or “unacceptable form of control” can be understood in some way broader than simply “a form of control that involves being at the mercy of decisions made on certain grounds.” If not, then it seems a counterexample to the Conjecture.

8. Cases: Discrimination

Another seeming counterexample appears in Scanlon's explanation of why racial discrimination—e.g., refusing to rent to a couple because they are of different “races”—is wrong. Part of his explanation is that this participates in and sustains a broader social practice that unfairly deprives members of a particular group of certain goods. Here, the reasons for which one acts do not matter; participating in the practice inadvertently would be wrong for the

same reason. The other part of Scanlon's explanation, however, is that this is a kind of insult, that it expresses a judgment of inferiority, which does seem to depend on the agent's reasons (p. 73). Scanlon suggests that racial insults are particularly significant when and because the judgment of inferiority is widespread. More idiosyncratic insults do not contribute to such prevalent patterns of judgment, and so they are less objectionable. Why this should mean, as Scanlon seems to take it to mean, that idiosyncratic insults are therefore permissible is less clear. In any event, he holds at least that when one's action is racially insulting, or expresses a judgment of racial inferiority, it is impermissible, and impermissible because of the reasons for which one acts. So this seems another counterexample to the Conjecture.

One might here try to defend the Conjecture with a maneuver similar to that proposed at the beginning of our discussion of incentive threats in the previous section. One might argue that it is not one's actual intent, but instead the intent that (one has reason to believe that) others have reason to attribute to one that matters for permissibility in this case. Suppose that in refusing to rent to the couple, one is not acting from racist motives at all, although there is every reason for them and others to misinterpret one's actions as having that meaning. (A stock scenario of a certain kind of contemporary comedy.) Would that make the refusal permissible (controlling for the considerations having to do with the deprivation of opportunities)? Not obviously. If the renter realized that his action was liable to be interpreted in that way, we might expect him to feel no longer morally free to proceed as he had planned, unable to continue justifying it to himself by thinking, "But that's not what I really mean." However, this is not an option that Scanlon himself takes up.

9. The phenomenology of deliberation

This brings us to the fourth sort of argument that might seem to support the Conjecture. In deliberating about what one may permissibly do, it seems somehow irrelevant to reflect on the reasons for which one would be acting. Scanlon illustrates this with the following example:

Suppose you were prime minister, and the commander of the air force described to you a planned air raid that would be expected to destroy a munitions plant and also kill a certain number of civilians, thereby probably undermining public support for the war. If he asked you whether you thought that this was morally permissible, you would not say, “Well, that depends on what your intentions would be in carrying it out. Would you be intending to kill the civilians, or would their deaths be merely an unintended but foreseeable (albeit beneficial) side effect of the destruction of the plant?” (pp. 19–20).¹⁶

In this particular case, it does seem odd for one’s deliberation to turn inward, to one’s own state of mind. However, this phenomenological point loses its force as soon as one tries to apply it to the firing and discrimination cases. It hardly seems odd to say to the employer: “Well, be honest with yourself. Are you firing her because her sales have dipped, or because she wouldn’t sleep with you?” Or to the owner renting the house: “Are you passing them over because they are a mixed-race couple?”

10. Re-interpreting the Conjecture?

In light of Scanlon’s endorsements of these apparent counterexamples to the Conjecture, is there another way to understand whatever general view on the relevance of intent to permissibility it is

¹⁶ Putting this speech in the mouth of (one presumes) the civilian commander-in-chief does make it especially odd, but in an irrelevant way. If the prime minister gives the go-ahead, then the raid will be, in an important sense, *his* doing. So what is he up to asking, or asking only, what the *air commander’s* intentions would be? If anyone’s intentions matter, don’t his own? There’s a curious abdication of responsibility.

that Scanlon really holds? Even in the firing case, Scanlon cautions us against a natural misconception about how intent matters:

What is morally basic is not a class of reasons that it is inherently illegitimate for agents to act on but, rather, the reasonable claims of victims not to have these agents' discretion to give or withhold goods defined so broadly as to make them unjustifiably subject to their control (86).

One might think at first that the misconception that Scanlon is guarding against is simply that intent matters nonderivatively: a point that might be put by saying, "It isn't the illegitimacy of the reasons that's really doing the work, but something more basic." But as we have seen, this clearly isn't what he means here.

One possible alternative, which is suggested by "not a class of reasons that it is inherently illegitimate for agents to act on," is that Scanlon's general view is:

There is no kind of reason such that *every* action for a reason of that kind is impermissible.

This view would allow a principle that forbids secretly firing an employee in order to get back at her for refusing one's advances, so long as the principle permitted *some other* action for that reason (such as *secretly sticking pins in a voodoo doll* or *saving the life of her rival* in order to get back at her for refusing one's advances). In other words, this view would rule out only the idea that the taint of a bad motive *suffices* to make an action impermissible.¹⁷

¹⁷ Believing is Necessary, May Implies Can Choose, and Always a Permissible Option would support this view. Let B be some bad reason, thought to make any action performed for it impermissible. Let X-ing be some action, which seems, at least when done for the right reason, required: e.g., saving many lives, or refraining from a grave injustice. Now imagine that one believes, albeit falsely and unjustifiably, that the only reason for X-ing is B. By Believing is Necessary, one cannot choose to X but not for B. So, by May Implies Can Choose, X-ing but not for B is not permissible. By Always a Permissible Option, if X-ing for B is indeed

Another possibility, which seems closer to the truth, and which is suggested by “but, rather, the reasonable claims of victims,” is that Scanlon’s general view is:

If the reason for which an action is performed makes it impermissible in a nonderivative way, it does so only in virtue of supporting some recognizable complaint or objection on the part of the victim.¹⁸

This view can accommodate the spiteful firing and discrimination cases, since there we can identify recognizable complaints: complaints about being under another’s control, or about being insulted, being judged to be inferior.

In principle, this general view might make harmless malice, such as the voodoo case, impermissible. It would turn on whether victims have recognizable complaints against the permissibility of such actions and how these complaints compare with the claims of agents to be allowed to engage in them. Indeed, once one approaches the problem from this more theoretical direction, and once one grants that there are valid complaints to others’ expressing the judgment that one is inferior, one wonders why harmless malice should be not only blameworthy, but also simply impermissible.¹⁹ Doesn’t one express the judgment that a person is inferior when one chooses to make that person the target of a plan that one believes, or ought to believe, is impermissible?²⁰ (It is true that this judgment may not be part of a widespread social pattern, as

impermissible, then not X-ing—e.g., not saving the lives, or committing the grave injustice—is permissible, which seems implausible.

¹⁸ Indeed, in comments, Scanlon suggests that this is the general view that he endorses. Compare Frances Kamm, “Terrorism and Intending Evil,” *Philosophy & Public Affairs* 36 (2008): 157–186, at pp. 168–69, who suggests that if an intention to harm or terrorize did make an action impermissible, it would have to do so in virtue of supporting some objection on the victims’ behalf.

¹⁹ We assume here there is some permissible, choosable alternative. Otherwise, we would have a helpful malice case.

²⁰ Taking another tack, might there not be a recognizable complaint against being made (by choice) the mental object of a violent, sexual, or degrading fantasy (without one’s consent): a

judgments of racial inferiority are. But as we mentioned earlier, it is not clear why this means that the target has no objection to the judgment, rather than simply that the target has a weaker objection.) What claim do agents have to being allowed to form such plans that might outweigh the legitimate objections of victims?

Something similar might even be proposed for the doctrine of double effect. In his defense of the doctrine, Warren Quinn suggests that what makes harm by “direct agency” especially objectionable is that the victim “fall[s] under [the agent’s] power and control in a distinctive way” (p. 348).²¹ Such “forced strategic subordination”²² of the victim to the agent’s ends, he concludes, is “especially inappropriate among free and equal” persons (p. 351). These considerations have at least the shape of complaints on the victim’s behalf. “Strategically subordinating” someone, as Quinn seems to view it, involves a kind of judgment that he is inferior: a mere tool. Couldn’t these complaints make Mixed Tactical Bombing impermissible: the case of moral opportunism, in which one bombs the munitions plant in order to demoralize the population by means of civilian deaths?²³ Wouldn’t this be a version of the traditional doctrine?

complaint against one’s mental image, so to speak, being used in such a way, as one might have about one’s photographic image being used in such a way? If so, it would be strange if one could not also have a complaint about being made the object of a violent, sexual, or degrading, albeit harmless, *plan*. In *What We Owe to Each Other* (Cambridge, Harvard: 1998), Scanlon seems at least to entertain this possibility: “[I]t might be said that masturbation generally involves fantasizing about other people, and that when it is morally objectionable this is because the attitudes it involves are incompatible with what we owe to others. This may be true in some cases...” (p. 175). I’m indebted here to a conversation with Jeremy Carey.

²¹ “Actions, Intentions, and Consequences: The Doctrine of Double Effect,” *Philosophy and Public Affairs* 18 (1989): 334–351.

²² Quinn’s conception of “subordination,” which depends on what the agent intends, should be distinguished from Kamm’s conception, which depends only on the non-psychological relation between the lesser evil and greater good that the action effects. See note 13.

²³ At least in a case in which there was a permissible, choosable alternative, such as not bombing at all or, if Believing is Sufficient is false, bombing for purely tactical objectives.

Of course, we may disagree with Quinn that such “subordination” is something that we have reason to object to, whether on the grounds that it expresses a judgment of inferiority or otherwise (see pp. 118–121).²⁴ At any rate, what seems to me right is that even if subordination would make Mixed Tactical Bombing impermissible, and even if subordination does count in some way against the permissibility of Terror Bombing, it simply does not seem important *enough* to begin to explain the difference between our reactions to Pure Tactical Bombing and Terror Bombing. Something like Scanlon’s alternative explanation must be what is really bearing the load.

These “concessions” may seem to make Scanlon’s view of the *general* relation between permissibility and intent too weak to be worth defending—setting aside the obvious value, independent of such a general view, of his distinction between permissibility and meaning, and his analyses of particular cases. To my mind, though, of the general views about the relation between permissibility and intent that we have considered, it is this view that most clearly bears on broader issues about the nature of morality. Why, after all, should it matter whether or not intent affects permissibility? One answer is that if intent affects permissibility, it may seem that this can be explained only by the idea, best suited to a religious or perfectionist doctrine, that the permissibility of an action depends on the purity of heart, or excellence of character, that an agent displays in performing it: on whether or not the action would be a good, or morally worthy, action. Scanlon touches on this possibility, when he observes that “supporters of the doctrine of double effect” may “believe that an agent, in deciding what to do, should be guided by whether a proposed action is good” (p. 24). On this view, one might say, the victim and what happens to

²⁴ Of course, *if* subordination is impermissible, then the thrust of the preceding discussion is that intending subordination *will* express a judgment of inferiority. But we are here looking for reasons why subordination might be impermissible in the first place.

him figure, in the first instance, only as temptations to vice, or prompts to virtue. This view is incompatible with Scanlon's account of what wrongness is: roughly, that what it is for an action to be impermissible is for those who would be affected by such actions to have a sufficiently compelling objection to principles allowing them. So, insofar as one finds Scanlon's account attractive, as I do, one has reason to resist this view. But one doesn't need to be a contractualist to think that this kind of inward focus, on whether one's motives are free of stain, distorts what matters for permissibility. The problem is not with the idea that intent, as such, should matter to permissibility, but rather with the idea that it should matter in *this* way: in virtue of what it says, in the first instance, about the *agent*.

If these are the reasons why it matters whether intent affects permissibility, then this last view about the general relation between permissibility and intent—that when intent affects permissibility, it does so by grounding some objection that those affected could raise—as weak as it may at first seem, is precisely the general view that matters.²⁵ And I suppose that this is where my mind comes to rest: *so long as* we can find convincing explanations in terms of valid objections of those affected, of why intent affects permissibility when it does, it is not clear why

²⁵ There is a parallel here to Scanlon's claim, made elsewhere and only touched on in *Moral Dimensions*, that the agent's aims do not give her reasons in general. As one probes this claim, one finds all sorts of exceptions. For example, the agent's friendships are partly constituted by his intentions, and yet they give him reasons. What seems to remain, after these exceptions are taken into account, is the claim that if intentions give us reasons, they do so in virtue of affecting some value, such as when friends' intentions *partly constitute* the value of friendship. And this residual claim may seem too weak to be of interest. Yet it is the version of the claim that most clearly bears on a broader issue about the nature of practical reason. Why, after all, should it matter whether or not intentions affect reasons? One answer is that it seems to challenge the idea that practical reasoning is a matter of responding to values. If intentions affect our reasons—the challenge would seem to go—this can be explained only by the idea that practical reasoning is a different kind of activity altogether: “self-legislation,” or the “creation” of reasons. If this is why it matters whether intentions affect reasons, then the residual claim is precisely what matters. To defend the idea that practical reasoning is a matter of responding to value, we need to show either that the intentions do not affect our reasons, or that, when they do, they do so only by affecting values. I discuss this at greater length in “Aims as Reasons.”

the idea that intent does affect permissibility—even, say, in the way that the doctrine of double effect predicts—should be puzzling or dismaying, as a matter either of philosophical theory or moral conviction.