



Moral Conflict and Its Structure

Author(s): David O. Brink

Source: *The Philosophical Review*, Apr., 1994, Vol. 103, No. 2 (Apr., 1994), pp. 215-247

Published by: Duke University Press on behalf of Philosophical Review

Stable URL: <https://www.jstor.org/stable/2185737>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



JSTOR

Duke University Press and are collaborating with JSTOR to digitize, preserve and extend access to *The Philosophical Review*

Moral Conflict and Its Structure

David O. Brink

It is commonly thought that there might be conflicts of obligations or moral dilemmas and that their existence would have serious implications for moral theory, metaethics, or deontic logic. Some alleged implications of moral dilemmas depend upon their *frequency*. On one such view, the fact that moral dilemmas are so common places severe constraints on the prospects for a systematic moral theory and on the possibility of moral knowledge.¹ But a more common view is that something about the *structure* of a moral dilemma generates troublesome or paradoxical implications. On this view, the important question is whether there are or could be any moral dilemmas; their frequency is of little or, at most, secondary importance. Indeed, if we understand 'moral dilemmas' univocally in both contexts, then the question of their existence must be prior to the question of their frequency. In any case, I shall focus on this second, structural issue.

Before we debate whether there are moral dilemmas, we need to know what sort of conflict is necessary to produce trouble or paradox in ethical theory. In fact, it will be useful to reserve the expression 'moral dilemma' as a term of art that refers to conflicts with such consequences. When the nature of these conflicts and their consequences are clear, we can ask whether it is reasonable to suppose that there are moral dilemmas.

1. *Prima Facie* and All-Things-Considered Obligations

A moral dilemma involves some kind of conflict of obligations or duties. (I shall not distinguish between obligations and duties. Moreover, I shall assume that both duties and obligations are ascribed using the moral "ought.") If an agent is faced with a moral

¹See, for example, Thomas Nagel, "The Fragmentation of Value" reprinted in his *Mortal Questions* (New York: Cambridge University Press, 1979) and Charles Taylor, "The Diversity of Goods," in *Utilitarianism and Beyond*, ed. A. Sen and B. Williams (New York: Cambridge University Press, 1982).

dilemma, he has an obligation to do A and an obligation to do B, but can't do both. This much is clear.² But we need to know more about the structure of these obligations. It is common and, I believe, important to distinguish between *prima facie* and *all-things-considered* obligations or duties. For, I shall argue, it is a conflict of all-things-considered, and not merely *prima facie*, duties that is required to generate a moral conflict that would have serious implications for ethical theory.

I propose to understand the distinction between *prima facie* and all-things-considered obligations in the way that I think W. D. Ross understood his distinction between *prima facie* and *sans phrase* obligations.³ A *prima facie* obligation to do x means that there is a moral reason to do x or that x possesses a right-making characteristic. But *prima facie* obligations can be, and often are, defeated by other, weightier obligations, individually or in concert. A *prima facie* obligation to do x that is superior to all others constitutes an all-things-considered obligation to do x. An all-things-considered moral obligation to do x means that on balance, or in view of all morally relevant factors, x is what one ought to do or that x is supported by the strongest moral reasons.

If *prima facie* obligations correspond to the presence of morally relevant factors or right-making characteristics, and an all-things-considered obligation is an undefeated *prima facie* obligation, then a natural way to understand a *prima facie* obligation to do x is as the claim that *ceteris paribus*, x is all-things-considered obligatory. The fact that x is *prima facie* obligatory results from x's possessing some morally relevant factor F and can be derived from the fact that x is F, together with the generalization that F-ness makes actions *prima facie* obligatory. This generalization should

²Notice three things about my discussion. First, I am discussing *moral* conflicts. There are comparable issues involving conflicting nonmoral oughts and conflicts between moral and nonmoral oughts, but these will not be my focus. Second, my focus is on conflicts *per accidens*, rather than conflicts *per se*. Most of my conclusions are easily carried over to conflicts *per se*. Third, I shall not distinguish between conflicts whose existence is the agent's own fault (*secundum quid*) and those that are not (*simpliciter*). These two sorts of conflicts are interestingly different. But these differences seem irrelevant to my main claims.

³See W. D. Ross, *The Right and the Good* (Oxford: Oxford University Press, Clarendon Press, 1930), 1–20, 28–29, and *Foundations of Ethics* (Oxford: Oxford University Press, Clarendon Press, 1939), 84–86.

itself be understood as the claim that *ceteris paribus*, F-ness makes actions all-things-considered obligatory. All else will be equal only if the moral factor F is *undefeated*. F will be undefeated if

- (a) there are no competing moral factors,
- (b) there are competing factors that cancel each other out, or
- (c) competing factors not canceled out do not override F's support for x.⁴

On this account, we treat *prima facie* obligations as moral factors or forces that interact so as to determine all-things-considered obligations. To determine all-things-considered obligations we must do *moral factor addition*. It is not essential to the factor addition model that we always be able to assign precise numerical values to the various moral forces present in a situation. What is important is that the moral status of an act *sans phrase* results from adding the moral forces, positive and negative, contributed by the various morally relevant factors; the act with the highest moral total is all-things-considered obligatory.⁵ This, I believe, is a fairly traditional account of *prima facie* and all-things-considered obligations that would have appealed to Ross.

⁴My views about the role of *ceteris paribus* clauses in laws and generalizations have been heavily influenced by Paul Pietroski, *Meaning Naturally: A Partial Defense of Covariation Semantics* (Ph.D. diss., MIT, 1990), esp. chaps. 4–5. Also see Paul Pietroski, “*Prima Facie Obligations, Ceteris Paribus Laws in Moral Theory*,” *Ethics* 103 (1993): 489–515.

⁵The factor addition model makes the *additive assumption* that the moral status of an act *sans phrase* results from adding the moral forces contributed by the various morally relevant factors. This may seem to involve the *atomist assumption* that the value or force of an individual factor is independent of the other factors in the situation. And the atomist assumption may be suspect. Perhaps certain factors (for example, pleasure) have positive moral force only when combined with other factors (for example, when they are innocent pleasures). If so, doubts about the atomist assumption may lead us to question the additive assumption and the factor addition model. See Shelly Kagan, “The Additive Fallacy,” *Ethics* 99 (1988): 5–31. But the additive assumption does not require the atomist assumption; we can reject the atomist assumption that the moral force of a factor is independent of all other factors and allow some moral factors to be specified relationally without giving up the additive assumption that moral status results from adding the values or forces of various moral factors. (Of course, we cannot specify moral factors in a *maximally* relational way if we are to preserve the distinction, essential to the factor addition model, between *prima facie* and all-things-considered duties or claims.)

Some writers think that *prima facie* obligations are not genuine obligations and that all genuine obligations are all-things-considered obligations.⁶ This tendency may be reinforced by the fact that the standard principles of deontic logic using the operator for obligation—‘O’—signify an all-things-considered obligation. This claim is harmless if it simply expresses our intention to call only all-things-considered moral requirements “duties” or “obligations” and to treat ‘*prima facie* obligation’ as a technical term. But I think that more than this is usually intended by those who deny that *prima facie* obligations are genuine obligations, and their denial rests on a misunderstanding of *prima facie* obligations that it is important to avoid.

These writers sometimes say that *prima facie* obligations are merely *apparent* obligations such that they have no moral force if overridden.⁷ But this does not fit our understanding of *prima facie* obligations or Ross’s. As Ross points out, we should not understand *prima facie* obligations as the *epistemic* claim that certain things appear to be obligatory that may not prove to be.⁸ This reading does not imply that there is any moral reason supporting *x* corresponding to the *prima facie* obligation to do *x*. Rather, *prima facie* obligations should be given a *metaphysical* reading that recognizes *prima facie* obligations as moral forces that are not canceled by the existence of other moral forces even if the latter override or defeat the former.⁹

Now Ross does say that *prima facie* duties are conditional duties

⁶Foot recognizes genuine obligations that may be overridden (type-1 obligations) and distinguishes them from the obligation associated with what there is the most moral reason to do (type-2 obligations), and so recognizes something like the distinction that I intend between *prima facie* and all-things-considered obligations. But she seems to treat *prima facie* obligations epistemically or statistically (see text below) and so does not want to equate the type-1/type-2 distinction with the *prima facie*/all-things-considered distinction. See Philippa Foot, “Moral Realism and Moral Dilemma,” reprinted in *Moral Dilemmas*, ed. C. Gowans (New York: Oxford University Press, 1987), 256–57. Because I reject these readings of *prima facie* obligations, our distinctions are similar.

⁷See Bernard Williams, “Ethical Consistency,” reprinted in *Moral Dilemmas*, ed. Gowans, 125, 126; Bas van Fraassen, “Values and the Heart’s Command,” *ibid.*, 141, 142; Ruth Barcan Marcus, “Moral Dilemmas and Consistency,” *ibid.*, 191; Foot, “Moral Realism and Moral Dilemma,” 257.

⁸*The Right and the Good*, 20.

⁹On the metaphysical reading, a *prima facie* obligation expresses a *pro tanto* moral obligation or moral reason.

and not duties proper.¹⁰ This, I believe, reflects only his decision to reserve the terms ‘duty’ and ‘obligation’ for all-things-considered moral claims. If we concede this to him, then we can explain most of his claims about prima facie obligations on our model. Prima facie obligations are *conditional* (all-things-considered) duties in the sense that if all else is equal, then there is not only a prima facie obligation to do x but also a genuine or all-things-considered obligation.

Sometimes Ross says that prima facie obligations refer to features of an act that *tend* to make acts of that type (all-things-considered) obligatory.¹¹ This claim admits of a *purely statistical* reading: though there may be nothing about this token act that makes it obligatory, it belongs to a type or class of acts many of whose tokens are (all-things-considered) obligatory.¹² But we can and should avoid this purely statistical reading of Ross’s tendency claims.

Suppose we have a prima facie obligation to do anything that is F and that we analyze this as the generalization that *ceteris paribus*, F-ness makes actions all-things-considered obligatory. Suppose also that all else is usually, but not always, equal. If so, the statistical claim is true. But it says nothing about F’s contribution to the rightness of right acts that are F or about F’s role in situations in which other things are not equal. But F does have a role and exerts a moral force in every situation in which it is present. In this way, the metaphysical reading of prima facie obligations explains why the statistical claim is true when it is true and implies reasonable claims that the statistical claim does not; so the statistical claim cannot exhaust our account of prima facie obligations.

Consider this parallel. We might analyze the claim that Bonny is a valuable player as the claim that she tends to help her team win games. But this should not be analyzed as the purely statistical claim that when she plays her team usually wins. This does not convey the idea that she helps *produce* their victories or the idea that she is a positive force even in the games they lose. We must understand her role as a positive factor when she’s on the field, whether the outcome is a win or a loss. This is how we must analyze

¹⁰*The Right and the Good*, 19.

¹¹*Ibid.*, 28–29.

¹²Cf. Foot, “Moral Realism and Moral Dilemma,” 257.

the tendency for her to win games if this claim is to be an analysis of her being a valuable player.

If so, we must reject the purely statistical reading. Indeed, Ross himself explains the tendency claims with an analogy between the way in which the interaction of prima facie obligations determines all-things-considered right conduct and the way in which vector addition determines the trajectories of physical objects in motion.¹³ This is precisely the “moral forces” interpretation of prima facie obligations.¹⁴

2. Soluble and Insoluble Conflicts and Regret

What does this account of prima facie and all-things-considered obligations have to do with moral dilemmas? If moral dilemmas are to be paradoxical things, then it’s arguable that the competing moral claims involved must be equipollent—neither claim should be stronger than the other. This is not just the epistemic point that the competing moral claims seem equally compelling; this would give us only an epistemic dilemma. Rather, genuine moral dilemmas must involve metaphysical equipollence; neither claim can be weightier than the other. If so, moral dilemmas must involve *insoluble* conflicts of prima facie obligations. For there seems nothing very puzzling about the existence of conflicting prima facie obligations one of which is stronger than the other.

However, not everyone thinks that moral dilemmas require insoluble conflicts. Consider a *soluble* conflict. There are moral rea-

¹³*The Right and the Good*, 29. The analogy is especially strong if moral factor addition *just is* moral vector addition. And perhaps it is. However, one apparent disanalogy is that in vector addition the output is rarely the same as any one of the inputs, whereas in moral factor addition the output is usually (always?) one of the inputs.

We might notice that Ross also connects prima facie obligations with *ceteris paribus* moral claims; see *The Right and the Good*, 30.

¹⁴We might also notice that our factor addition model of the relation between prima facie obligations and all-things-considered obligations does not imply consequentialism or utilitarianism. Though factor addition requires that we add the moral forces, positive and negative, associated with alternative actions in order to determine which action available to the agent is supported by the strongest moral force, this does not require assuming that the strength of a moral force is proportional to the amount of good or value that an action would realize. Thus, Ross can help himself to the factor addition model, and he is not thereby a closet utilitarian.

sons to do A and there are moral reasons to do B, but one cannot do both A and B. Moreover, we assume, though the reasons supporting A and those supporting B are both strong, there is in fact more reason to do A. Furthermore, to simplify, let's assume that the agent in question recognizes the greater stringency of A. Our traditional model implies that soluble moral conflicts have a perfectly straightforward analysis. What we have is one *prima facie* obligation being overridden or defeated by another; this weightier *prima facie* obligation is the agent's all-things-considered obligation. But some friends of dilemmas think that there is a genuine dilemma even in such a case provided that the less weighty moral claim grounds compunction or regret on the agent's part when she fails to act on it in order to fulfill her weightier obligation. They take the appropriateness of attitudes such as regret or compunction towards B to show that the moral reasons supporting B do not disappear just because the reasons for B are defeated by the reasons for A.¹⁵ And this must be incompatible with the traditional account, for such an account, as Williams claims, must "eliminate from the scene the *ought* that is not acted upon" and so remove the ground of any compunction or regret.¹⁶

Let us assume, for the sake of argument, that a kind of regret towards the fact of an unperformed obligation is appropriate.¹⁷

¹⁵See Williams, "Ethical Consistency," 122–25, 134; van Fraassen, "Values and the Heart's Command," 147–48, 151; and Marcus, "Moral Dilemmas and Consistency," 193, 196–97.

¹⁶"Ethical Consistency," 122–25.

¹⁷Is regret appropriate in soluble conflicts? First, we must distinguish the issue of whether it's appropriate to *harbor* or *express* regret from the issue of whether the thing in question is an *appropriate object* of regret. It is only the latter issue that is in question. Second, whether compunction or regret is appropriate depends, I think, on the *object* of such attitudes. We might distinguish some apparently different objects of regret. Regret might focus on (i) the circumstances that make some sort of conflict (now) unavoidable, (ii) the consequences for others of the unperformed obligation, (iii) the fact that one obligation remains unperformed, or (iv) the fact that one performed the action one did. (Cf. Foot, "Moral Realism and Moral Dilemma," 257.) Though (i)-type and (ii)-type regret are arguably appropriate, they do not threaten the traditional account; (iv)-type regret seems clearly inappropriate; only (iii)-type regret is potentially troublesome for the traditional account. Let me explain briefly.

The circumstances that make conflict unavoidable seem an appropriate object of regret. We wish the world was not such as to make it necessary to leave one of our *prima facie* obligations unfulfilled. But there also seems

The friends of dilemmas might have a reasonable argument against the traditional analysis of soluble conflicts if the traditional account relied on an epistemic or purely statistical reading of *prima facie* obligations. For then the weaker *prima facie* claim would or need be no real moral factor at all; indeed, on the epistemic reading, moral factors operate only when they win.¹⁸ If the reasons supporting B do not operate where B is overridden by A, then regret or compunction could apparently be appropriate only for *instrumental* reasons. If the agent did not feel compunction about her failure to perform B (where this caused someone's suffering), perhaps she would become more callous and fail to do B-type acts even when they are (all-things-considered) obligatory.¹⁹ This in-

to be no basis for paradox here and nothing that the factor addition model of soluble dilemmas cannot explain. Because we want to respond appropriately to all moral forces, we regret that circumstances require us to ignore some moral forces in order to attend to stronger ones.

(ii)-type regret is more complicated. Indeed, one may even wonder if this is a distinct category of regret. Typically, though not always, failure to meet a *prima facie* obligation results in hardship to the person(s) to whom the obligation was owed. We may regret the fact that this person must suffer this hardship and the fact that we play a role in producing this hardship. But this regret seems perfectly compatible with recognizing that it is morally necessary to cause this hardship in order to avoid violating stronger moral requirements. Perhaps this regret just involves the wish that things could have been otherwise such that causing this hardship was not necessary. If so, it involves no more than (i)-type regret.

The object of one's regret may be the unfulfilled, but weaker, obligation—(iii)-type regret. One regrets not performing B. But it's also not clear that this is a distinct kind of regret. One's reason for regret at not performing B might be the fact that one would have liked to perform B. But this is regret that one was in a situation such that performing B was morally unavailable (because it would have been morally irresponsible to act on the weaker claim). This is (i)-type regret again, and poses no problem to the traditional analysis of soluble conflicts. Alternatively, this regret might involve more; it might involve the thought that in the actual circumstances one should have performed B, rather than A. This is (iv)-type regret. But it seems precisely inappropriate. Surely, if it is a soluble conflict in which there is a greater obligation to do A, then it *is* inappropriate to wish that one had performed B, rather than A.

However, perhaps (iii)-type regret need not collapse into either (i)-type or (iv)-type regret. It may be appropriate to experience regret or compunction at the fact of not fulfilling an obligation, even if it is an obligation that is overridden by others. I discuss this form of regret in the text.

¹⁸Winning presupposes a contest, and a contest presupposes competitors. If so, the idea that a factor is a factor only if it wins may be incoherent.

¹⁹This may be Hare's view about the sort of regret that is appropriate

strumental account of the appropriateness of compunction or regret may not seem to ring true.

But if we accept the metaphysical reading of prima facie obligations and view the function from prima facie obligations into all-things-considered obligation as moral factor addition, then we can explain regret for failing to perform defeated prima facie obligations. If prima facie obligations are moral forces that are at work even when they are overridden by competing forces, then the fact that an agent does not act on her prima facie obligation to do B explains the appropriateness of regret or compunction. Insofar as B is prima facie obligatory, it should be performed; if I do not perform it, whether A overrides B or not, then I am not responding to moral forces present in the situation. And this might be thought to be a cause for genuine regret and compunction. Perhaps a morally decent person should have moral attitudes, such as compunction or regret, that track moral forces or factors. If so, I can acknowledge that regret or compunction is appropriate towards my failure to do B while admitting that it is A that is all-things-considered obligatory. These reactions are sufficiently accounted for, on our metaphysical reading, by the assumption that I had a prima facie obligation to do B.²⁰

3. A Recipe for Moral Dilemmas

For these reasons, we should deny that soluble conflicts are genuine moral dilemmas. Moral dilemmas must be insoluble, involving claims that are (metaphysically) equipollent. However, we need require only *broad equipollence*. It is not necessary that the competing claims represent the same denomination in a single moral currency. Between some currencies there may be no uniform exchange rate, with the result that we must recognize pockets of incommensurability. If two moral claims are incommensurably stringent, I assume that neither is greater or more stringent than the other. If so, the two claims are broadly equipollent.

If a genuine dilemma cannot be a mere conflict of prima facie obligations but must involve competing claims that are broadly

in cases of moral conflict; see R. M. Hare, *Moral Thinking* (New York: Oxford University Press, 1981), chaps. 2–3.

²⁰Cf. Ross, *The Right and the Good*, 28.

equipollent, we might construe moral dilemmas as competing all-things-considered obligations. If so, the recipe for moral dilemmas appears to be something like this.

1. One has a *prima facie* obligation to do A.
2. One has a *prima facie* obligation to do B.
3. One is under an all-things-considered obligation to do x just in case one is under a *prima facie* obligation to do x, and there is no greater, simple or complex, competing *prima facie* obligation one is under.
4. One's *prima facie* obligation to do A is no greater than one's *prima facie* obligation to do B, and vice versa.
5. One is under no other *prima facie* obligation, simple or complex, that competes with A or B and that is as great an obligation.
6. Hence one has an all-things-considered obligation to do A. [1,3–5]
7. Hence one has an all-things-considered obligation to do B. [2–5]
8. It is possible for one to do A.
9. It is possible for one to do B.
10. It is not possible for one to do A and B.

I trust that the recipe itself is reasonably clear. For now I just want to comment on the steps in the recipe that are less familiar and say enough to make them initially plausible. The less familiar steps are (3)–(5). (3) simply states the apparently traditional idea that an all-things-considered obligation is an undefeated *prima facie* obligation. (4)–(5) describe the conditions that must be met if a *prima facie* obligation is to be undefeated.

(4) may seem to be the crucial step in the recipe. Once we distinguish the epistemic question of knowing or having justified belief that a *prima facie* obligation is defeated and the metaphysical question of its being defeated, we may wonder whether there are conflicting undefeated *prima facie* obligations. We may admit that it is often hard to decide which obligation is weightier but insist that the epistemic difficulty in settling conflicts is no evidence that they're insoluble.

Our belief in the frequency of insoluble conflicts is likely to

depend upon what we make of the possibility of pockets of incommensurability involving moral claims of different kinds (for example, fidelity and nonmaleficence). Incommensurability presumably will exist only if there are different scales or dimensions of assessment. Let us say that there is *strong incommensurability* between types of acts x and y just in case no token of type x is comparable with any token of type y; and let us say that there is *weak incommensurability* between types of acts x and y just in case some tokens of type x are not comparable with some tokens of type y, but some are. Strong incommensurability is obviously more radical than weak incommensurability; weak incommensurability allows for substantial, though partial, comparability.²¹ Even if we have different dimensions of assessment, strong incommensurability seems implausibly extreme.²² However, weak incommensurability strikes me as fairly plausible. But while I'm sympathetic to the suggestion that pockets of incommensurability of this sort exist, I don't know how to argue for this claim and won't try.

A more secure route for the friend of dilemmas to take is to embrace the theoretical possibility of ties. The idea, as Ruth Marcus notes, would be to describe cases of symmetrically structured conflicting claims. We might imagine a parent who is equidistant between identical twins who are trapped in a burning building and can save one but not both. Make the case maximally symmetrical. The parent's duties conflict and their stringency is the same. Perhaps the choice that the Nazis force Sophie to make about which of her children to sacrifice involves such a conflict.²³ Narrow equi-

²¹Where partial comparability applies, the relation *not worse than* is not transitive. Cf. Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press, Clarendon Press, 1984), 431. Suppose that A and C are items on one dimension of assessment and B is an item on another and that A is marginally greater than C. If both A and C are incommensurable with respect to B, then C is no worse than B, and B is no worse than A. But if A is greater than C, C is worse than A; thus, the relation *not worse than* is not transitive.

²²A very strong version of incommensurability would be needed to show that insoluble moral conflict is so pervasive as to threaten the possibility of systematic moral theory and moral knowledge. It is in part because I see no motivation for this kind of incommensurability that I think it more profitable to examine the structural issues about moral dilemmas that are independent of issues about their frequency.

²³William Styron, *Sophie's Choice* (New York: Random House, 1979). Of course, Sophie's choice is an insoluble conflict only if it is false that her

pollence is not required, but it's more secure, because its theoretical possibility is harder to dispute.²⁴

4. Paradoxical Implications of Genuine Dilemmas

If we follow the recipe, then we have a conflict of all-things-considered obligations. We can then attach an obligation operator to both A and B. The reason for thinking that this is the structure of a genuine moral dilemma is that when we conjoin such a conflict with other initially plausible claims we get some very troublesome results. In fact, the existence of genuine moral dilemmas and the truth of these auxiliary claims generate contradictions.

In explaining these paradoxes, it will help to represent them symbolically. I adopt the following conventions. 'O(A)' means that one ought all-things-considered to do A or that A is all-things-considered obligatory. Indeed, unless otherwise indicated, I shall understand 'obligation' to refer to all-things-considered obligation. Where I have occasion to represent prima facie obligations symbolically I shall use the lower case 'o'; 'o(A)' means that one ought prima facie to do A or that A is prima facie obligatory. 'P(A)' means that A is permissible. '♦(A)' means that A is in the relevant sense possible or within the agent's power—roughly,

duty not to be involved in the sacrifice of any of her children is greater than her duty to save one of her children.

²⁴In this way I want to disagree with McConnell (see Terrance McConnell, "Moral Dilemmas and Consistency in Ethics," reprinted in *Moral Dilemmas*, ed. Gowans). He wants to reject moral dilemmas by arguing that once we distinguish epistemic and metaphysical issues, we should deny that there are or could be conflicts between undefeated prima facie obligations. But, while I agree that we should be skeptical about the frequency with which insoluble conflicts occur, I think that we must admit that they are possible, and the case of symmetrical but conflicting duties demonstrates this most clearly. I too will reject the possibility of moral dilemmas (see section 5), but because I reject (3), not because I believe insoluble conflicts are impossible.

Foot construes her type-2 oughts as claims about what there is most or best reason to do and concludes that there cannot be conflicts of type-2 oughts (see Foot, "Moral Realism and Moral Dilemma," 256). But conflicting actions can both be best as long as the corresponding moral claims are broadly equipollent. If so, then Foot has no ground for denying type-2 conflicts.

physically or psychologically possible.²⁵ For instance, in the example described above, it is not possible in the relevant sense, given familiar temporal and physical constraints, for the parent to save both of her twins from the burning building. Deontic principles or axioms (axiom schemata) are formulated using Greek letters (for example, ‘ α ’ and ‘ β ’). These principles do figure as axioms or theorems in various systems of deontic logic. However, I am concerned only with whether they are intuitively compelling claims about the systematic relations among familiar moral and nonmoral properties (for example, the relations between permissibility and obligatoriness). If a principle is intuitively compelling, this is presumably (defeasible) reason to think that it should figure as an axiom or theorem in any reasonable system of deontic logic.

Paradox 1

One paradox has been widely recognized. It results from combining moral dilemmas with two standard deontic principles: a distribution principle, often known as *agglomeration*, and the principle that “ought” implies “can,” which I’ll dub the *voluntarist* principle. Agglomeration says that if one ought to do α and one ought to do β then one ought to do both.

$$(O(\alpha) \ \& \ O(\beta)) \rightarrow O(\alpha \ \& \ \beta)$$

The voluntarist principle says that it is a necessary condition of

²⁵Thus, we should distinguish the modality that ‘ \blacklozenge ’ signifies from the sort of metaphysical possibility that ‘ \Diamond ’ signifies. Moreover, insofar as we are concerned with impossibility, it is with specific kinds of impossibility, not with the sort of impossibility that, if determinism is true, characterizes an agent’s relation to all actions she does not perform.

one's being under an obligation to do α that α be possible or be within one's power.²⁶

$$O(\alpha) \rightarrow \blacklozenge(\alpha)$$

By combining agglomeration and voluntarism we get the following paradox.

1. $O(A)$
2. $O(B)$
3. $\neg\blacklozenge(A \ \& \ B)$
4. $(O(A) \ \& \ O(B)) \rightarrow O(A \ \& \ B)$ [agglomeration]
5. Hence $O(A \ \& \ B)$ [1,2,4]
6. $O(A \ \& \ B) \rightarrow \blacklozenge(A \ \& \ B)$ [voluntarism]
7. Hence $\blacklozenge(A \ \& \ B)$ [5,6]
8. Hence $\blacklozenge(A \ \& \ B) \ \& \ \neg\blacklozenge(A \ \& \ B)$ [3,7]

How impressive are the credentials of the auxiliary principles used to derive this paradox? A full discussion is not possible here; and, because the other paradoxes are more robust, a full discussion is

²⁶Common sense distinguishes between two kinds of modalities. Most of us assume that alternative actions can each be within my power even if it is causally determined which one I will perform. But some actions that are possible in this quite general way are nonetheless impossible or outside an agent's powers because the particular context imposes a sufficient physical or psychological disability (for example, as where the parent is unable to save both twins from the fire). These actions are impossible for the agent in question in a way different from the way in which, if determinism is true, all the actions the agent does not perform are impossible. In formulating voluntarism and construing a moral conflict as a conflict among duties such that one can perform each but not both, I am relying on a conception of possibility that includes both the more general and the more specific forms of possibility. So if determinism and compatibilism are both true, A and B will each be possible in both the general and specific senses, but it will be impossible in the specific sense to do both. In such a case, the voluntarist can accommodate obligations to do each, but not an obligation to do both. If incompatibilism and determinism are true, then there is no distinction to draw between the modality in which it is impossible for the parent to save both children and the modality in which it is impossible to act otherwise than one does. A voluntarist would conclude that there can be obligations to do only things one does. But we would face no moral conflicts; A and B are not each within the agent's power, and, according to voluntarism, there could not be two conflicting obligations.

not necessary. It is enough to explore some central worries and resources.

Consider agglomeration. Where there is no conflict between A and B, it seems harmless to recognize an obligation to do both as well as obligations to do each. But it also seems unnecessary; an obligation to do each seems adequate to explain the moral situation.

However, perhaps I rely on agglomeration in *avoiding* dilemmas. Assuming that I want to fulfill my obligations, I want to avoid doing things that will prevent me from doing so. If I believe that I ought to do A and that I ought to do B, not only I will try to avoid doing things that will prevent me from doing one or the other, I will try to avoid doing things that prevent me from doing both. Does this presuppose that I have an obligation to do both as well as obligations to do each? I don't think so. All it requires is that I have two obligations and a desire to satisfy each; it does not require a conjunctive obligation.

Now, I can think of one argument for compound obligations. Consider a case where the agent has an obligation to do A and an obligation to do B and *is* able to do both A and B. In such a case it seems reasonable to hold that it is impermissible for him to fail to do both.

$$\neg P(\neg (A \& B))$$

But it also seems reasonable to accept a *correlativity* principle according to which α is obligatory just in case it is impermissible not to do α .

$$\neg P(\neg \alpha) \equiv O(\alpha)$$

But these two claims imply that he has a compound obligation.

1. $\neg P(\neg (A \& B))$
2. $\neg P(\neg (A \& B)) \equiv O(A \& B)$ [correlativity]
3. Hence $O(A \& B)$

This argument establishes that there are conjunctive obligations

where each conjunct is obligatory *and* both obligations can be performed. However, it does not establish, as agglomeration claims, that conjunctive obligations follow from the obligatoriness of each of the two conjuncts. Moreover, this particular route to conjunctive obligations won't work in the case of moral dilemmas. For, as I shall argue (section 5), in a genuine dilemma, where one cannot do both A and B, what is impermissible is to perform neither; it is not impermissible to fail to perform both, provided one performs one of them.

$$\neg\neg P(\neg(A \& B))$$

If so, the correlativity principle does not allow us to generate a conjunctive obligation in dilemmatic situations.

Indeed, we might think that it is the correlativity principle that explains why there is a conjunctive obligation when there is one. If so, the existence of conjunctive obligations in nondilemmatic situations provides no support whatsoever for agglomeration or its claim that there are conjunctive obligations in dilemmas.

Nor does there seem to be any reason to assume the existence of conjunctive obligations in dilemmatic situations, as agglomeration would require. Why should we assume that there's an obligation to do both? Why not say, instead, that there is an obligation to do A, an obligation to do B, *and* an obligation to do A or B, but *not* an obligation to do both?

What about voluntarism? It seems more intuitively compelling. It may seem too harsh to require of people actions that they cannot perform; the voluntarist insists that it be up to us whether to be moral. Nonetheless some people claim that there are counterexamples to the voluntarist principle; they claim that there are impossible obligations. We will have more confidence in discarding the voluntarist principle if there are counterexamples in nondilemmatic situations.

Some of the counterexamples involve obligations that one is no longer in a position to fulfill. For example, it has been claimed that a debtor who is no longer able to pay his debt, whether through his own fault or not, is still under the obligation to repay his debt.²⁷

²⁷See, for example, Michael Stocker, "Moral Conflicts: What They Are and What They Show," *Pacific Philosophical Quarterly* 68 (1987): 104–23.

But it seems that we can admit that the debtor remains under obligations to those whom he owes without assuming that he has an obligation to perform the repayment that he is incapable of performing. If the debtor is himself at fault for becoming unable to pay the debt at the appointed time, then when he fails to pay the debt he can be blamed, not for failing to pay the debt, but for allowing himself to become unable to pay the debt. Moreover, whether his inability to pay is his own fault or not, we can claim that *ceteris paribus*, he has an obligation to pay back as much as he can until the original debt is paid, perhaps with interest or compensatory damages. These claims seem not to require positing impossible obligations.

However, it might be claimed that duties of restitution or compensation make sense only as a consequence of breaching an obligation. If so, a duty to compensate seems to imply that the debtor was obligated to do something he was unable to do. But this analysis is not mandatory. Some obligations get met; others lapse or end without being discharged. And obligations can end or lapse for a variety of different reasons. Some of the reasons an obligation lapses may bring no new obligations in its wake, as when a creditor chooses to release the debtor from his obligation. But other reasons for the lapse of an obligation may bring new obligations in its wake. If a debtor is unable to pay as agreed, then his old obligation lapses and, *ceteris paribus*, is succeeded by new obligations to repay according to a new schedule and to compensate for damages. The voluntarist can appeal to this view about the *succession* of obligations to recognize duties of compensation without recognizing impossible obligations. If so, it's not clear that we yet have a counterexample to the voluntarist principle.

Another possible challenge to voluntarism involves moral emotions. People often say that someone ought to have some moral feeling or attitude (for example, gratitude) in a particular situation even if the person does not have the attitude and seems incapable in the circumstances of producing it. While emotional sensibilities may be intentionally cultivated or repressed, it seems they are not under our direct voluntary control in the required sense. If we have not cultivated or otherwise acquired the relevant background sensibilities, it may not always be within our power to experience the appropriate attitudes in the appropriate contexts. Yet it may nonetheless seem true to say that we ought to have these attitudes.

But we may question the counterexample without disputing the psychological assumptions upon which it relies. For one thing, it's not clear that this is a case of an impossible obligation. While it may seem natural to say that the person in question ought to have these feelings, we need not be ascribing an obligation to him. We often say that something ought to be the case without ascribing obligations to anyone; here we usually mean that it would be a good thing if what ought to happen did, or perhaps that it would have been a good thing if what ought to have happened had. No one thinks that voluntarism must apply to these axiological uses of 'ought'. Moreover, we may think that the person in question has or had certain obligations, even if he does not have an impossible obligation. For if he does not have sensibilities that allow him to have the appropriate feelings in this situation, he may well be blameworthy for having violated an obligation, which was in his power to perform, to cultivate the relevant sensibilities and he may now have an obligation (to do what he can) to cultivate the relevant sensibilities. So, it's not clear that there really is any obligation in such cases that the voluntarist cannot recognize.

I haven't shown voluntarism to be immune to counterexample (and there are the usual problems about proving negative existential generalizations). But my analysis of these counterexamples suggests voluntarist strategies for dealing with other alleged counterexamples. To say that someone ought to have done something she was unable to do may reflect an axiological, rather than a deontic, judgment; if so, we can accept the judgment without recognizing impossible obligations. Moreover, we may recognize that the agent in such a situation has or had obligations without recognizing impossible obligations. There may have been an obligation the agent violated in getting into the predicament, and there may be obligations to mitigate damages and provide restitution. We need to test the adequacy of these strategies on a case-by-case basis.

Paradox 2

The second paradox also relies on two deontic principles. The first says that if you're (all-things-considered) obligated to do α , and β would prevent α or bring about not- α , then you're obligated not to do β .

$$(O(\alpha) \ \& \ (\beta \rightarrow \neg \alpha)) \rightarrow O(\neg \beta)$$

We might call this the *obligation execution* principle, because it obligates us not to do anything that would interfere with the execution of our (original) obligations.

The second principle simply claims that if it is obligatory to do not- β , then it is not obligatory to do β .

$$O(\neg \beta) \rightarrow \neg O(\beta)$$

This statement of the consequences of an obligation is weaker than a third principle we are likely to find attractive.

$$O(\neg \beta) \equiv \neg P(\beta)$$

The third principle asserts a kind of *correlativity* of obligatoriness and impermissibility that is equivalent to the correlativity principle we've already discussed ($\neg P(\neg \alpha) \equiv O(\alpha)$).²⁸ This version of correlativity says that if not- β is obligatory, then β is impermissible, and vice versa. Because the second principle is weaker than this correlativity principle (and for lack of a better name), I shall call it the *weak obligation* principle. The weak obligation and obligation execution principles help constitute a second paradox of moral dilemmas.

1. $O(A)$
2. $O(B)$
3. $\neg \blacklozenge (A \ \& \ B)$
4. Hence $B \rightarrow \neg A$ [3]
5. $(O(A) \ \& \ (B \rightarrow \neg A)) \rightarrow O(\neg B)$ [obligation execution]
6. Hence $O(\neg B)$ [1,4,5]
7. $O(\neg B) \rightarrow \neg O(B)$ [weak obligation]
8. Hence $\neg O(B)$ [6,7]
9. Hence $O(B) \ \& \ \neg O(B)$ [2,8]

²⁸The two correlativity principles are equivalent as long as we adopt a rule of extensionality allowing us to substitute logical equivalents within the scope of the deontic operators 'O' and 'P'. Then, if we treat either principle as an axiom schema, we can derive the other as an instance.

Now, the basis of the second paradox seems much more secure than that of the first paradox.

The obligation execution principle surely plays an important, though perhaps tacit, role in our moral deliberations. Consider some examples. I recognize an obligation not to kill innocent people; I know or learn that putting a certain substance (in an unlabeled spice bottle) in Alan's food would bring about his death; so I conclude that I ought not to put the substance in his food. Or, I recognize an obligation to be with my son on his birthday; I receive an invitation to participate in a conference on that day; so I conclude that I ought to refuse this invitation. Or, I recognize an obligation to meet you at the time and place we agreed upon; I realize that the only way for me to do so is to catch the noon train; so I conclude that I ought to catch the noon train. These common inferences seem reasonable, because obligatoriness seems to be transmitted from the action that is the focus of the original obligation to those actions that are required, as means or necessary conditions, for the performance of the original obligation.²⁹ It's

²⁹Similar principles are discussed by van Fraassen, "Value and the Heart's Command," 146, and by Judith Thomson, *The Realm of Rights* (Cambridge: Harvard University Press, 1990), 156–57.

It should be clear that this requires us to read the embedded conditional in the obligation execution principle as asserting some kind of counterfactual or nomological dependence between β and $\neg\alpha$, rather than as a simple material conditional. For instance, we might represent the embedded conditional with something like Stalnaker's "corner"; see Robert Stalnaker, "A Theory of Conditionals," reprinted in *Causation and Conditionals*, ed. E. Sosa (New York: Oxford University Press, 1975).

Someone might think obligation execution too strong a principle. By contraposition within obligation execution we get $(O(\alpha) \ \& \ (\alpha \rightarrow \neg\beta)) \rightarrow O(\neg\beta)$. But it may seem that some nomic consequences (for example, epiphenomena) of things we are obligated to do need not themselves be obligatory. For instance, suppose I am obligated to give a talk on a certain date at a certain university, one consequence of which is that I see something of the campus after the talk and before returning home. It might seem strange to suppose that I am obligated to see the campus as well as give my talk. Upstream consequences of our obligations may be obligatory, but downstream consequences do not seem to be. But I'm not so sure. If downstream consequences are genuine consequences, then failure to bring them about entails that one will not have fulfilled one's original obligation. *Ceteris paribus*, if I do not see the campus, then I will not have fulfilled my obligation to give my talk. So in one sense failure to realize downstream consequences of our obligations prevents our fulfilling those obligations much as failure to realize upstream consequences does. If so, perhaps both

hard to imagine how we would find our way around in the world morally if we did not accept this principle.

And the weak obligation principle, as its name suggests, seems especially uncontroversial. If I'm obligated not to kill my neighbor, then surely it's not the case that I'm obligated to kill him. Indeed, as the stronger, correlativity principle asserts, if I'm obligated not to kill my neighbor, then it's impermissible for me to kill him. But our second principle is even weaker than this plausible third principle; so surely it must be acceptable.

Paradox 3

A third paradox results from combining the obligation execution and correlativity principles with a principle that claims that if β is impermissible then it is not the case that β is obligatory.

$$\neg P(\beta) \rightarrow \neg O(\beta)$$

Because this is a very weak principle about the consequences of impermissibility, I will call it the *weak impermissibility* principle. These three principles and the assumption of moral dilemmas create the following paradox.

1. $O(A)$
2. $O(B)$
3. $\neg \Diamond(A \ \& \ B)$
4. Hence $B \rightarrow \neg A$ [3]
5. $(O(A) \ \& \ (B \rightarrow \neg A)) \rightarrow O(\neg B)$ [obligation execution]

sorts of consequences are obligatory and obligation execution is defensible. Moreover, the friend of obligation execution can recognize an important *pragmatic* asymmetry between upstream and downstream consequences. Upstream consequences are or can be part of planning the execution of one's original obligations, as downstream consequences cannot.

Alternatively, of course, we might defend obligation execution simply by interpreting the embedded conditional as representing a narrower kind of nomological dependence between an obligation and its upstream consequences of the sort when β would prevent α from coming about.

- | | |
|--------------------------------------|-------------------------|
| 6. Hence $O(\neg B)$ | [1,4,5] |
| 7. $O(\neg B) \equiv \neg P(B)$ | [correlativity] |
| 8. Hence $\neg P(B)$ | [6,7] |
| 9. $\neg P(B) \rightarrow \neg O(B)$ | [weak impermissibility] |
| 10. Hence $\neg O(B)$ | [8,9] |
| 11. Hence $O(B) \ \& \ \neg O(B)$ | [2,10] |

We might also notice that the obligation execution and correlativity principles allow us to generalize the argument for (8) to the claim that in a dilemma, whatever the agent does, she does something impermissible. For these two principles imply that any action that leads to the nonperformance of an obligation is impermissible, and the impossibility of doing both A and B ensures that whatever the agent does she will fail to perform an obligation. It follows that whatever she does, she does something impermissible. We might consider this result overly harsh; it is incompatible with the general voluntarist idea that it ought to be within our power to live up to our obligations and to avoid impermissible actions. In this sense, some may regard this result as a distinct paradox of dilemmas. However, this conclusion does not itself involve a contradiction and so is not paradoxical in the strong sense that the other results are.

Nonetheless, (11) is a contradiction. As with the second paradox, the auxiliary deontic principles in the third paradox seem especially compelling. We've already seen the plausibility of the obligation execution and correlativity theses. The new deontic principle is the weak impermissibility principle. But surely that must be true. If it's impermissible for me to torture my neighbor, then surely it's not the case that I'm obligated to torture him.

5. Reconsidering the Existence of Moral Dilemmas

Because the paradoxes are generated from the supposition that there are moral dilemmas, understood as conflicts of all-things-considered obligation, and various deontic principles, we must reject some of the deontic principles or the supposition that there are moral dilemmas. Which reaction is correct depends on the independent plausibility of the principles and the case for dilemmas. We must determine which reaction is least costly. The principles underlying the first paradox—agglomeration and volunta-

rism—strike me as somewhat uncertain. Rejecting one or both of them—especially agglomeration—may be a cost that we could bear. However, the second and third paradoxes rest on more secure deontic foundations. Obligation execution, weak obligation, correlativity, and weak impermissibility all seem to be fundamental or uncontroversial principles in our moral reasoning. If so, the second and third paradoxes make the costs of maintaining the existence of moral dilemmas greater than has generally been recognized.

We could simply restrict the scope of such principles to nondilemmatic situations in order to avoid paradox. But that would be ad hoc. We need a principled and independent account of why their scope should be limited in just this way. Moreover, if we restrict the scope of these deontic principles to nondilemmatic situations, then different principles will apply to obligations in dilemmatic and nondilemmatic contexts. But if the sense of deontic terms (such as ‘obligation’) is determined jointly by the various deontic principles in which the terms figure, then this response to the paradoxes appears to have the unwelcome consequence of requiring us to say that the term ‘obligation’ has different meanings in the two contexts. If so, it is better to reject such principles than to restrict their scope. But then we must find alternative mechanisms (of unrestricted scope) for generating their results in non-dilemmatic situations. However, I see no such mechanisms to serve as alternatives to the principles underlying the second and third paradoxes.³⁰ If I am right, friends of dilemmas reject such principles at a high price.

These paradoxes suggest that we should examine the case for moral dilemmas more closely. Earlier, I presented a *prima facie* case for thinking that moral dilemmas are theoretically possible. And they need be only theoretically possible—people need never face a genuine dilemma—in order for ethical theory to contain paradox. Though I agreed that we should be skeptical about the frequency of conflicting undefeated *prima facie* obligations, I ar-

³⁰In this connection, it is worth noting that my view about agglomeration is not that we should just restrict its scope to nondilemmatic situations. That would be ad hoc. I think that we should reject it and that we can identify a *different* mechanism for generating compound obligations in non-dilemmatic situations, namely, the first correlativity principle. However, I do not see any similarly plausible alternatives to obligation execution, correlativity, weak obligation, or weak impermissibility.

gued that we must allow for their possibility. They are possible either because of weak incommensurability between prima facie obligations or because of equal stringency of prima facie obligations, such as cases of symmetrically structured, conflicting obligations. As long as we accept the idea that an undefeated prima facie obligation yields an all-things-considered obligation, we must admit that moral dilemmas are possible.

It is this idea that we can and should reject. Ordinarily, an undefeated prima facie obligation does constitute an all-things-considered obligation. But not always. Where there is an undefeated competitor, we can conclude that neither obligation is an all-things-considered obligation. This may seem to leave the agent confronting an insoluble conflict with no all-things-considered obligations, and this may seem puzzling to some. But the agent does face an all-things-considered obligation; it is to perform one or the other of the conflicting prima facie obligations. We can see this better if we allow 'O' to stand for all-things-considered obligation (as usual) and 'o' to stand for prima facie obligation. In an insoluble conflict of undefeated prima facie obligations, the following claims seem true.

- o(A)
- o(B)
- $\neg (o(A) > o(B))$
- $\neg (o(B) > o(A))$
- $O(A \vee B)$
- $\neg O(A)$
- $\neg O(B)$

If so, the only all-things-considered obligation in an insoluble conflict is this disjunctive obligation.³¹ And the disjunctive all-things-considered obligation feels right. By correlativity, if the agent performs neither disjunct, she has done something impermissible. But she may perform either of the disjuncts in order to fulfill her all-things-considered obligation. Whichever disjunct she performs, she will leave a strong (indeed, undefeated) prima facie obligation un-

³¹Alan Donagan discusses some aspects of the disjunctive solution in "Consistency in Rationalist Moral Systems," reprinted in *Moral Dilemmas*, ed. Gowans, 286–87.

performed, and this may be cause for a kind of regret or compunction for the moral force to which she does not respond. But as long as she performs one of the disjuncts, she will have done nothing impermissible.

This disjunctive analysis seems fairly natural as applied to narrowly equipollent moral conflicts; in the case of a tie, it seems permissible to perform either alternative. But the analysis can and should be extended to moral conflicts whose insolubility is due to incommensurability. For what makes the disjunctive analysis appropriate is just that when all the moral considerations about the alternatives have been exhausted neither moral claim is weightier than the other. And this is true for any insoluble conflict, whether the insolubility represents a tie or incommensurability.

If, as I believe, the disjunctive analysis of insoluble conflicts is perfectly adequate, then we should deny the possibility of moral dilemmas.³² In particular, we should reject step (3) in the recipe

³²This form of skepticism about moral dilemmas should sound familiar, because it is similar to my criticism of agglomeration. Here we deny the existence of genuine dilemmas by admitting only a disjunctive all-things-considered obligation. There we argued that if there were moral dilemmas, we might avoid the first paradox by denying agglomeration, and we might do this by conceding only a disjunctive obligation. We admit that an (all-things-considered) obligation to do A and an obligation to do B yields a disjunctive obligation to do A or B, but we deny that it yields a conjunctive obligation to do both A and B. But the present point is not simply the denial of agglomeration for *prima facie* obligations. We can see this symbolically. When we deny agglomeration we assert four main claims.

- O(A)
- O(B)
- O(A \vee B)
- \neg O(A & B)

If we denied the “agglomeration” of *prima facie* obligations we would assert four parallel claims.

- o(A)
- o(B)
- o(A \vee B)
- \neg o(A & B)

Instead, in denying the possibility of dilemmas, we make the following

for moral dilemmas and its claim that an all-things-considered obligation is simply an undefeated *prima facie* obligation. An all-things-considered obligation represents what one ought to do in light of *all* morally relevant factors, including alternatives. If so, then only *prima facie* obligations that are undefeated *and* defeat all competitors are all-things-considered obligations. In other words, to be an all-things-considered obligation, a *prima facie* obligation must be *overriding* and not simply not overridden. But there cannot be two incompatible obligations of this sort; it cannot be true that A overrides B and that B overrides A.

Recognition that insoluble conflicts generate only a disjunctive obligation itself provides a kind of solution to these conflicts. There is no all-things-considered obligation to do A or B, rather than the other. But there is an all-things-considered obligation to do one or the other, rather than some third thing. In particular, there is an obligation to do one or the other, rather than nothing. To fulfill this disjunctive obligation one must choose one or the other, though, by hypothesis, the choice cannot be guided by moral considerations about A or B. Instead, one's choice should be determined by the greater salience of one disjunct or by some (other) randomizing device. In this way, our analysis does not itself make the choice. But it does tell us that we have an unequivocal duty to do either A or B and that, as far as our duty goes, it doesn't matter which we do. This is itself a kind of guidance.

This parallels what I take to be the right response to another sort of insoluble conflict—the case of Buridan's Ass. Neither bale of hay is more desirable than the other, and so there is no reason for the Ass to prefer one bale rather than the other. Nonetheless

claims about insoluble conflicts.

- $o(A)$
- $o(B)$
- $O(A \vee B)$
- $\neg O(A \& B)$
- $\neg O(A)$
- $\neg O(B)$

The explanation of these differences is that whereas agglomeration is a function that takes only all-things-considered obligations as arguments, premise (3) in our recipe for a dilemma takes *prima facie* obligations as input and all-things-considered obligations as output. So, though similar in part, skepticism about dilemmas and skepticism about agglomeration have different structures.

there is a very strong reason for the Ass to choose one of the bales, it doesn't matter which, rather than starve in indecision. The rational Ass, therefore, should quit looking for reasons for choosing one, grounded in the desirability of the bales, and choose the more salient bale or employ a randomizing device to select one. Only in this way will he do what is rationally required of him, namely, to eat from one or the other bale. We give the Ass a kind of guidance when we tell him it doesn't matter which bale he chooses so long as he chooses one.

However, this disjunctive solution to insoluble moral conflicts may seem to trivialize them. It is commonly thought that such conflicts present No Win situations, whereas the disjunctive solution may seem to treat such conflicts as No Lose situations. More generally, the disjunctive solution may seem to take the conflict out of moral conflict.

Now it is true that the disjunctive solution implies that in an insoluble conflict an agent does nothing impermissible as long as she performs either A or B, it doesn't matter which.³³ But, whatever she does, she will fail to perform a strong—indeed, undefeated—prima facie obligation, and this can be an appropriate object of compunction or regret. So the disjunctive solution allows that whatever the agent does, it may be appropriate for her to experience compunction or regret. In this way, the disjunctive solution can represent insoluble conflicts as No Win situations. Moreover, the disjunctive solution does not imply that it's easy to recognize genuinely insoluble conflicts. Indeed, if the broad equipollence of competing moral claims is rare enough and hard enough to identify reliably, it may never be reasonable for an agent to conclude that she is in an insoluble conflict to which the disjunctive solution applies (even if she is in one).³⁴ So the disjunctive solution preserves robust epistemic conflict for the agent in an insoluble con-

³³To be accurate, the only insoluble conflicts that the disjunctive solution must represent as No Lose situations are conflicts *per se*, in which A and B are mutually exclusive and jointly exhaustive alternatives; in such a case, whatever one does, one performs one of the disjuncts in the disjunctive obligation. Conflicts *per accidens*, in which A and B are mutually exclusive but not jointly exhaustive, are not No Lose situations; one does something impermissible if one performs neither disjunct.

³⁴For a rigorous statement of a similar claim about the judicial treatment of legal conflict, see Rolf Sartorius, "Bayes' Theorem, Hard Cases, and Judicial Discretion," *Georgia Law Review* 11 (1977): 1269–75.

flict. If we bear these forms of residue in mind, I see no reason to find the disjunctive solution implausibly sanitary.

6. Moral Conflict and Metaethics

Williams believes that moral dilemmas have metaethical implications as well as implications for deontic logic.

It seems to me a fundamental criticism of many ethical theories that their accounts of moral conflict and its resolution do not do justice to the facts of regret and related considerations: basically because they eliminate from the scene the *ought* not acted upon. A structure appropriate to conflicts of beliefs is projected onto the moral case; one by which the conflict is basically adventitious, and a resolution of it disembarasses one of a mistaken view which for a while confused the situation. Such an approach must be inherent in purely cognitive accounts of the matter, since it is just a question of which of the conflicting *ought* statements is true, and they cannot both be true, to decide correctly for one of them must be to be rid of error with respect to the other.³⁵

Moral conflicts, he concludes, “are more like conflicts of desires than they are like conflicts of beliefs.”³⁶

Insofar as Williams’s claim is simply that a cognitivist construal of conflicts of obligation cannot accommodate the phenomena of regret or compunction, his argument has already been addressed (in section 2). As long as we treat *prima facie* obligations as genuine moral factors or forces in the situations to which they apply, we can treat soluble conflicts as cases in which one *prima facie* obligation overrides another, and insoluble conflicts as cases of conflicting undefeated *prima facie* obligations. In either case, we can explain such regret as is appropriate for the *prima facie* obligation not acted upon.

But Williams also seems to be offering a somewhat different argument against cognitivism. He wants to compare conflicts of beliefs and conflicts of desire or attitude, and argue that moral conflicts have properties more like those of conflicts of attitude than those of conflicts of belief. His argument seems to have the following form.

³⁵“Ethical Consistency,” 125.

³⁶*Ibid.*, 121.

1. It is unreasonable to hold conflicting (contradictory) beliefs.
2. It can be reasonable to hold conflicting noncognitive attitudes towards the same thing.
3. In moral conflicts, it is reasonable for the agent to make conflicting judgments about her obligations.
4. Hence judgments and ascriptions of obligation in cases of moral conflict express attitudes, rather than beliefs.
5. There should be a uniform semantic treatment of judgments and ascriptions of obligation.
6. Hence judgments and ascriptions of obligation express attitudes, rather than beliefs.

Williams thinks that (2) is obvious, but I'm skeptical. It can be reasonable to hold different attitudes towards different aspects of the same situation. But it's less clear that it can be reasonable to hold different attitudes towards one and the same aspect of one and the same situation. There's no problem taking conflicting attitudes (for example, attraction and repulsion) towards an object if we take one attitude toward it qua F and take the other attitude toward it qua G. For then our attitudes take different objects—they are directed at *different aspects* of one and the same object—though at a superficial, coarse-grained level of description they both take x as their object. Thus, I can be both pleased and dismayed by my toddler's new resistance to being dressed: I am pleased insofar as this means that he is developing a sense of his own independence (qua F) but dismayed insofar as dressing him is now less pleasant and more time-consuming for both of us (qua G). Indeed, it is these different properties of the same event that *ground* the appropriateness of my conflicting reactions, presumably because our attitudes towards objects are based on their properties. If so, this is not an example of having conflicting attitudes towards one and the same thing.

It's much less clear that it can be reasonable to hold conflicting attitudes towards x qua F—where this can't be explained away as having different attitudes towards different constituent properties of the complex property F. If it is aspects of a situation that ground or make appropriate certain attitudes towards it, then there is no reason to expect that it could be reasonable to hold conflicting

attitudes towards one and the same aspect of a situation.³⁷ At least, not unless we've *already* been convinced of a noncognitivist thesis to the effect that one and the same natural property of a situation can make different responses appropriate.

What about (3)? Is it reasonable for the agent to make conflicting judgments about her obligations? If (3) is to support the non-cognitivist conclusion, the conflicting judgments must endorse contradictory propositions about her obligations.

If my analysis is correct, it cannot be reasonable to accept contradictory propositions about the agent's all-things-considered obligations. There are no conflicting all-things-considered obligations. In soluble conflicts, there is only one all-things-considered obligation, and that is to perform the weightier *prima facie* obligation. In insoluble conflicts, the only all-things-considered obligation is the disjunctive obligation.

However, there can be conflicting *prima facie* obligations (in both soluble and insoluble conflicts). Do these conflicts yield contradictory propositions about the agent's *prima facie* obligations? A and B are each *prima facie* obligatory. Presumably we can accept an analogue of obligation execution for *prima facie* obligation.

$$o(\alpha) \ \& \ (\beta \rightarrow \neg \alpha) \rightarrow o(\neg \beta)$$

Because A and B are incompatible, my *prima facie* obligation to do either (for example, A) generates a *prima facie* obligation not to do the other (for example, B). This allows one to argue as follows.

³⁷In a similar way, Greenspan appeals to emotional conflict to argue against a cognitive construal of the emotions; see Patricia Greenspan, "A Case of Mixed Feelings: Ambivalence and the Logic of Emotion," in *Explaining Emotions*, ed. A. Rorty (Los Angeles: University of California Press, 1980). Though I do not accept the view that she calls cognitivism—the claim that emotions just are judgments or beliefs—I do think that this and other cognitivist views of the emotions are untouched by her argument. In particular, a cognitivist view that holds that emotions are and should be grounded in beliefs about the natural and evaluative properties of situations can readily accommodate conflicts of attitudes, because, as I indicated, these conflicting attitudes take different aspects of the same situation as their objects. The parallel between Greenspan's and Williams's arguments occurred to me during discussions with Tracy Isaacs.

1. $o(A)$
2. $o(B)$
3. $\neg \Diamond (A \ \& \ B)$
4. Hence $B \rightarrow \neg A$ [3]
5. $(o(A) \ \& \ (B \rightarrow \neg A)) \rightarrow o(\neg B)$ [“obligation execution”]
6. Hence $o(\neg B)$ [1,4,5]
7. Hence $o(B) \ \& \ o(\neg B)$ [2,6]

But, of course, this conclusion falls short of the required contradiction

$$o(B) \ \& \ \neg o(B).$$

Nor do I see how Williams can get the negation outside of the scope of the *prima facie* obligation operator. My having reason not to perform an act does not prevent there being reason to perform it. Indeed, more generally, analogues of weak obligation, correlativity, and weak impermissibility all fail when applied to *prima facie*, rather than all-things-considered, obligation. So premise (3) of Williams’s argument is false.

This should be clear from our account of *prima facie* obligations. Recall that *prima facie* obligations represent the operation of moral forces or factors and that we can represent these moral forces as moral laws or generalizations. These generalizations imply that an action is obligatory *insofar* as it possesses a morally relevant property. In moral conflicts, one and the same action has more than one morally relevant property. Action A is obligatory *insofar* as performing it responds to one moral force, while it is not obligatory *insofar* as performing it fails to respond to another moral force. There is no contradiction here, because the logical form of our predications is this:

A is O *insofar* as it is F

and

A is not-O *insofar* as it is G.

Williams's problem is that he sees contradictory attitudes and obligations where there are none, and this is because he fails to see that the attitudes and obligations in question attach to different aspects of the same situation. Because of this fact, we can see that there is no contradiction. So we should reject this argument for noncognitivism based on moral conflict.

7. Conclusion

There may well be insoluble moral conflicts, perhaps a great many. I have not addressed directly the frequency of such conflicts or the significance of their frequency.³⁸ But moral conflict of this sort does not itself have paradoxical implications for ethical theory, as friends of moral dilemmas have claimed.

Whereas the robust nature of the second and third paradoxes *motivates* a reexamination of the case for moral dilemmas, the argument against moral dilemmas is largely *independent* of the deontic paradoxes and the plausibility of the deontic principles. An all-things-considered moral obligation is what one ought to do in light of all morally relevant factors. To be an all-things-considered obligation, a *prima facie* obligation must not only be undefeated, it must defeat competitors; it must be overriding. And, indeed, familiar deontic principles, of the sort discussed here, are plausible only when applied to all-things-considered obligations, understood in this way. But there cannot be conflicting overriding requirements. So there cannot be conflicting all-things-considered obligations. In insoluble conflicts, neither of the conflicting obligations is an all-things-considered obligation; the only all-things-considered obligation is the disjunctive obligation to perform one or the other of the conflicting undefeated *prima facie* obligations. If we under-

³⁸Issues about the frequency of moral dilemmas need not be resolved by our conclusion that there can be no moral dilemmas, provided we do not understand 'moral dilemmas' univocally in both contexts. Whereas the structural debate concerns the existence of conflicting all-things-considered obligations, and I have argued that there can be none, the frequency debate is easily understood as concerning the existence and frequency of insoluble conflicts between conflicting undefeated *prima facie* obligations. My own view about this, for which I have not argued here, is that there can be and are such conflicts but that there is no good reason to think them so common as to be troublesome for ethical theory.

stand moral dilemmas as conflicts of all-things-considered obligations, to which familiar deontic principles apply, as many writers have, then we should agree that moral dilemmas would produce paradox in ethical theory but deny that there are or could be any moral dilemmas.³⁹

We can deny the possibility of moral dilemmas without disputing any of the moral phenomenology to which friends of moral dilemmas point. Indeed, both the construction and the resolution of the paradoxes require us to understand better the distinction between *prima facie* and all-things-considered obligations and to do justice to *prima facie* obligations. Alleged dilemmas are really conflicts between undefeated *prima facie* obligations. Such conflicts can be unfortunate but are not paradoxical. Indeed, when we understand *prima facie* obligations as genuine moral factors or forces in situations, we will better understand why insoluble and even soluble conflicts can be unfortunate and even tragic.⁴⁰

University of California, San Diego

³⁹Sinnott-Armstrong represents himself as a friend of dilemmas; he construes them as conflicts among non-overridden moral requirements and denies that they have paradoxical implications for ethical theory. See Walter Sinnott-Armstrong, *Moral Dilemmas* (Oxford: Blackwell, 1988). Because non-overridden moral requirements just are undefeated *prima facie* obligations, I agree that *if* we understood moral dilemmas as involving such conflicts, then moral dilemmas would have no paradoxical implications. But Sinnott-Armstrong must deny, what other participants in the debates (for example, Williams, van Fraassen, Marcus, Foot, and McConnell) assume, that moral dilemmas involve conflicts of all-things-considered obligations, to which familiar deontic principles apply.

⁴⁰This paper was written during a fellowship at the Center for Advanced Study in the Behavioral Sciences that was funded by an Old Dominion Fellowship from the Massachusetts Institute of Technology and by grants from the National Endowment for the Humanities (#RA-20037-88) and the Andrew W. Mellon Foundation. I would like to thank these institutions for their support. I would also like to thank Norman Dahl, Alan Donagan, Catherine Elgin, Ted Everett, Mark Johnston, Terrance McConnell, Alison McIntyre, Ruth Marcus, Gene Mason, Paul Pietroski, Peter Railton, Amélie Rorty, Alan Sidelle, Walter Sinnott-Armstrong, Robert Stalnaker, Peter Vranas, Susan Wolf, the editors of the *Philosophical Review*, and audiences at the University of Minnesota Conference on Moral Dilemmas, Princeton University, the University of Vermont, Brown University, the University of Wisconsin at Madison, the University of Arizona, Tufts University, the University of Virginia, the University of Southern California, and the University of California at San Diego for helpful comments and discussion.