

Goodness and Advice

JUDITH JARVIS THOMSON

THE TANNER LECTURES ON HUMAN VALUES

Delivered at

Princeton University
March 24 and 25, 1999

JUDITH JARVIS THOMSON is professor of philosophy at the Massachusetts Institute of Technology. She was educated at Barnard College and at Cambridge University, and she received her Ph.D. from Columbia University. She is a fellow of the American Academy of Arts and Sciences and has also been a fellow at the Centre for Advanced Study, Oslo, Norway. She has been a visiting fellow or professor at the Australian National University, the University of Pittsburgh, the University of California at Berkeley, and Yale Law School. She is a member of the American Philosophical Association and is the recipient of a Guggenheim Fellowship. She has published in numerous philosophical journals, including *Philosophy and Public Affairs*, the *Journal of Philosophy*, and *Nous*, is the author of *Acts and Other Events* (1977); *Rights, Restitution, and Risk* (1986); *The Realm of Rights* (1990); and co-author of *Moral Relativism and Moral Objectivity* (1996, with Gilbert Harman).

PART ONE: GOODNESS

I.

Twentieth-century Anglo-American moral philosophy has been dominated by concern about the fact-value gap. Or at least about what appears to be a gap, indeed, an unbridgeable gap, between fact and value. Matters of fact seem to be epistemologically intelligible: we find out about them by the familiar methods of observation and experiment. Matters of value seem to be quite different. If we can't learn about them by reasoning to them from matters of fact, then there seems to be no way at all by which we can come to learn about them. But what reasoning could possibly take a person from a matter of fact to a matter of value? It is hard to see how any reasoning could. Are we therefore to conclude that nobody has good reason to believe about any judgment of value that it is true? Many moral philosophers regard that as an appalling conclusion, and try to show that it is unwarranted. Others think it the correct conclusion, and try to show that we should not be troubled by it. In any case, all start from the apparent fact-value gap; responding to the threat it seems to pose became the central task of Anglo-American moral philosophy in the century just past.

That there does not merely seem to be, that there really is a fact-value gap, is by now part of the culture. Or at least, part of a certain culture, that of the middle-class literate public. I don't need to *introduce* my freshmen to the fact-value gap: they bring it to college with them.

I do not mean that members of the public at large, and my college freshmen, do not have moral beliefs. Their acceptance of the idea that there is a fact-value gap shows itself in more subtle ways. If I ask my students whether it is morally permissible for people to lie and cheat whenever it would profit them to do so, they reply, "No, of course not."

Delivered as a Tanner Lecture on Human Values at Princeton University, 1999. What follows is a revised and expanded version of the lectures. I am grateful to my four commentators, and to the participants in the discussions following the lectures, for very helpful comments and criticism. Some of their comments have been taken into consideration in revising the text.

Material in Part I, sections 9–14, was part of a paper presented in a symposium with Dennis Stampe at the meeting of the American Philosophical Division, Central Division, in May 1999. I thank Stampe and the other participants for their comments and criticism.

If I go on to ask what they would say if a man disagreed, they do not reply, “He’s mistaken, for the following reason: . . .”—giving a reason. What emerges is rather: “It’s all just a matter of opinion. I have mine and he has his, and they differ.” Or: “It’s all just a matter of how you feel. I feel one way, he feels another.”

We might describe the situation in this way. People have a great many beliefs about what it is good or bad, right or wrong to do; these are their first-order moral beliefs. What many people lack nowadays is the second-order belief that they have good reason to believe that their first-order moral beliefs are true.

I suppose that in some people, lack of this second-order belief shows itself in lack of confidence in their first-order moral beliefs. I doubt that this is true of many people, however. For the most part, I think, people who think that there is an unbridgeable fact-value gap, and therefore conclude that nobody has good reason to believe about any value judgment that it is true, feel entirely confident that lying and cheating are wrong. In most people, I think, lack of the second-order belief shows itself only on occasions on which they step back from ordinary life and reflect on their first-order moral beliefs—as, for example, in classrooms that say “Philosophy” on the door.

So where’s the harm in it? I said just above that many moral philosophers regard the conclusion that nobody has good reason to believe about any judgment of value that it is true as appalling. Why so?—if accepting that conclusion is compatible with feeling entirely confident that lying and cheating are wrong.

That these are compatible is something that moral philosophers who accept the conclusion try to demonstrate. Perhaps they are right. Presumably it is not literally inconsistent in a man to feel entirely confident that a certain number will win tomorrow’s lottery while believing that he has no good reason to believe it will. At best, however, he has a divided consciousness. And it would plainly be silly in him to stake much on that number.

Do those who accept the conclusion that nobody has good reason to believe about any judgment of value that it is true, while nevertheless feeling entirely confident that lying and cheating are wrong, have a divided consciousness? And would it be silly in them to stake much on its being the case that lying and cheating are wrong? These are good questions. But the prior question is surely whether anyone should accept that conclusion.

We should go back further, in fact. *Is* there really an unbridgeable gap between fact and value?

The prevalence of the idea that there is such a gap must have deep sources outside philosophy. That there is such a gap is a philosophical thesis, and no philosophical thesis becomes part of the culture unless there are other ideas in the culture that it reinforces and is reinforced by.

The prevalence of the idea that there is an unbridgeable gap between fact and value is overdetermined; in addition to sources outside philosophy, it has sources in other areas of philosophy as well as within moral philosophy. What I refer to are the grounds for a quite general skepticism that were laid out so strikingly by Descartes. My freshmen bring that too to college with them: they say that of course they believe they have fingers and toes, but they also say “That’s just my opinion” and “That’s just how I feel.”

What I will fix on in Part I is one of the considerations within moral philosophy itself that has led moral philosophers to regard the idea as at a minimum having to be taken seriously. I will suggest that the consideration I discuss is a product of illegitimate abstraction.

2.

Since well before the twentieth century, moral philosophers have taken it to be their task to produce a theory about what we ought to do and why. That “why” is important: moralizers are happy to tell you what you ought to do—moral philosophers differ in that they aim to tell you also what makes it the case that you ought to do the things they say you ought to do. Moral philosophy, in other words, responds to the desire that moral requirement be ‘rationalized,’ that is, shown to be a requirement.

Very well: what might be thought to make it the case that Alfred in particular, or people generally, ought to do this or that?

Suppose that Alfred acted in a certain way—he took off his hat, or pressed a certain doorbell, or what you will. Let us say he pressed a doorbell. For him to do that was for an event to occur, namely, the event that consisted in his pressing the doorbell.

We may suppose that Alfred’s pressing the doorbell caused many other events to occur. Thus, his pressing the doorbell caused a circuit to close, a bell to ring, a person inside the house to feel pleased, and so on

and on. These events—Alfred’s pressing the doorbell, the circuit’s closing, the bell’s ringing, and so on—are all of them events that would not have occurred if Alfred had not pressed the doorbell.

More generally, for a person to act is for a battery of events to occur, events that would not have occurred if he had not acted. We could put it more grandly: for a person to act is for the world to go in a way that it otherwise would not.

If that is what it is for a person to act, then it is very natural to think that the question whether he ought to act in such and such a way must turn on a comparison between what the world will be like if he acts in that way and what the world will be like if he acts in any of the other ways in which it is open to him to act. So, for example, that the question whether Alfred ought to press the doorbell turns on a comparison between what the world will be like if he presses it and what the world will be like if he instead stamps his feet, or dances a jig, or stands stock still, or There seems to be nothing else for it to turn on.

And what are the terms of the required comparison? An answer that all but suggests itself is this: we need to settle whether what the world will be like if he acts in such and such a way is *better or worse than* what the world will be like if he acts in any of the other ways in which it is open to him to act. Thus if the world will be better if Alfred presses the doorbell than it will be if he instead stamps his feet, and better than it will be if he instead dances a jig, and so on, then he ought to press the doorbell. And if the world will not be better if Alfred presses the doorbell than it will be if he instead does one of those other things, then it is not the case that he ought to press the doorbell.

In sum, a person ought to do a thing if and only if the world will be better if he does it than if he does any of the other things it is open to him to do at the time. Similarly, a person ought not do a thing if and only if the world will be worse if he does it than if he does any of the other things it is open to him to do at the time.

What if there is a tie? For example, what if the world will be equally good whether Alfred presses the doorbell or dances a jig, but better if he does either than if he does anything else it is open to him to do? The idea we are looking at yields that it is not the case that Alfred ought to press the doorbell and that it is not the case that he ought to dance a jig (since the world will not be better if he does either). But the idea also yields that it is not the case that Alfred ought not press the doorbell, and that it is not the case that Alfred ought not dance a jig (since the world will not be worse if he does either). So the idea yields that he need not, but

may, do either. It also yields, however, that he ought not do anything other than either.¹ These conclusions are plausible enough, and the possibility of ties therefore does not constitute a difficulty for those who are attracted by this idea.

Following current usage, I will call the idea Consequentialism.² It is, I think, deeply satisfying. How *could* it be perfectly all right to do a thing if the world will be worse if you do it than if you do something else instead? Moreover, given that for a person to act just is for the world to go in a way that it otherwise would not go, surely the question whether he ought to act had better turn on a comparison between how it will go if he acts and how it will go if he does something else—to repeat, there seems to be nothing else for it to turn on. And how is that comparison to be made if not by settling on which way of going would be better? Down the road from this idea lies the fact-value gap, among other serious difficulties. I have wanted first to bring out the idea's attractions. It is not surprising that so many people have found it attractive, and that those who reject it do not merely dismiss it, but feel the need to make a case against it.

I stress that Consequentialism says nothing at all about what would make the world be better or worse than it otherwise would be. The idea itself leaves that open.

But a moral philosopher needs to arrive at a view about this, so let us turn to it.

3.

What would make it the case that the world will be better if a person does one thing than it will be if he does another? “Better” is just the comparative of “good,” so we can re-put our question as follows: what

¹These are G. E. Moore's conclusions about ties; see his *Ethics* (London: Oxford University Press, 1949), 22–25. I thank Martha Nussbaum for reminding me of the need to mention ties.

²The idea was first given this name by G. E. M. Anscombe, in “Modern Moral Philosophy,” *Philosophy* 33 (1958). Her article was reprinted in G. E. M. Anscombe, *Collected Philosophical Papers*, vol. 3 (Oxford: Basil Blackwell, 1981). The name is unfortunate, and more's the pity that it has become common usage, since it is so very likely to mislead. A Consequentialist does not believe that what fixes whether a person ought to do a thing is a comparison between the consequences of his doing it with the consequences of his doing anything else, if “consequences” is understood in the most natural way, namely as “effects.” A Consequentialist believes that what fixes whether a person ought to do a thing is rather a comparison between what the world will be like if he does it with what the world will be like if he does anything else. These are very different ideas. For more on the difference, see note 3 below.

would make it the case that the world will be more good if a person does one thing than it will be if he does another? It is very natural to think that the world will be more good if it contains more of what is good or less of what is bad or both. Consider Alfred again. I invited you to suppose that if Alfred presses the doorbell, a great many events will occur that otherwise would not, namely, his pressing the doorbell, the circuit's closing, the bell's ringing, someone's feeling pleased, and so on. We may similarly suppose that if Alfred instead dances a jig, a great many events will occur that otherwise would not. Will the world be better, that is, more good, if he presses the doorbell than if he dances a jig? That—it is very natural to think—turns on whether some or other of the events that will occur if he presses the doorbell, and some that will occur if he instead dances a jig, will be good or bad, and if so, on how good or bad they will be.

Which events are good and which bad? A familiar idea says that an event is good just in case it consists in someone's feeling pleased, and bad just in case it consists in someone's feeling pain. This idea comes down to us from Bentham and John Stuart Mill.

It is certainly possible for Jones to be more pleased than Smith is. Suppose that is now the case. If an event is good if and only if it consists in someone's feeling pleased, then presumably the event that consists in Jones's being pleased is better than the event that consists in Smith's being pleased. Similarly for pains: if Jones's pain is more severe than Smith's, then the event that consists in Jones's feeling pain is worse than the event that consists in Smith's feeling pain.

I will call this idea about which events are good or bad, and about how good or bad they are, Hedonism About Goodness. Many people have found it a very attractive idea.

To return to Alfred, then. If he presses the doorbell, a battery of events will occur. If he instead dances a jig, a different battery of events will occur. Hedonism About Goodness tells us which of those events are good and which bad, and how good or bad they are. I mentioned that if Alfred presses the doorbell, a person inside the house will be pleased. Suppose that if Alfred instead dances a jig, then Alfred will be pleased. Suppose that Alfred will be less pleased if he dances a jig than the person inside the house will be if Alfred presses the doorbell. That counts in favor of its being the case that the world will be better if he presses the doorbell than if he dances a jig. But of course we would need to know a good deal more if we were to arrive at a conclusion on this matter: we

would need to know whether anyone else will be pleased if Alfred acts in each way, and moreover, whether anyone will feel pain if he does, taking into consideration everyone who would be affected by Alfred's pressing the doorbell and by Alfred's instead dancing a jig.

Still, if we conjoin Hedonism About Goodness with Consequentialism, we have produced a theory about what a person ought to do and why. A person ought to do a thing if and only if—and if so, *because*—the balance of pleasure and pain that ensues if he does it is greater than that which ensues if he does any of the other things it is open to him to do instead.

The idea we have reached is, of course, Utilitarianism. Utilitarianism is one version of Consequentialism: a Consequentialist is also a Utilitarian if and only if he accepts Hedonism About Goodness.³ I have laid out the process of reasoning by which it may be reached at some length, in order to bring out that it relies on two ideas, which it is important to distinguish from each other. One is Consequentialism, the other is Hedonism About Goodness. The contemporary literature of moral philosophy is full of objections to Utilitarianism; it is important, however, to be clear which of those two ideas a given objection to Utilitarianism is an objection to.

I am going to focus on an objection to Consequentialism that I think has not been taken seriously enough. But let us begin with Hedonism About Goodness.

I tried to make Hedonism About Goodness seem plausible when I presented it, which from this vantage point—the end of the twentieth century—is not easy to do. I know of no moral philosophers nowadays

³I said in note 2 above that the following two ideas are very different: (i) what fixes whether a person ought to do a thing is a comparison between the effects of his doing it with the effects of his doing anything else, and (ii) what fixes whether a person ought to do a thing is a comparison between what the world will be like if he does it with what the world will be like if he does anything else. And I said that it is (ii), not (i), that a Consequentialist opts for.

The difference between those ideas emerges clearly only if a Consequentialist rejects Hedonism About Goodness. Suppose a Consequentialist accepts Hedonism About Goodness. Then on his view, what matters to the question whether Alfred ought (as it might be) to press the doorbell is only the effects of Alfred's doing so and the effects of Alfred's acting otherwise. For example, while if Alfred presses the doorbell, then the world will contain the event that consists in his pressing the doorbell, that is morally irrelevant, since the event that consists in his pressing the doorbell is not itself, but at most has among its effects, events that consists in someone's being pleased or feeling pain.

But a Consequentialist might reject Hedonism About Goodness. In particular, he might instead say that some acts are themselves good or bad. On his view, then, it is not only events that are the effects of an act that matter morally, and opting for (ii) may yield a moral conclusion that is different from the moral outcome yielded by (i).

who accept it. Among many other objections, it has often been pointed out that a man might be pleased at someone else's feeling pain. Is his feeling pleased really to be thought a good event? We are surely inclined to think it positively vicious in a person to take pleasure in the pains of others. Insofar as we have intuitions about what counts as a good event—and I will return to this caveat shortly—it strikes us, intuitively, that a man's feeling pleased at the pain of another is not a good event.

Moreover, there is an interplay between Hedonism About Goodness and Consequentialism. Suppose we accept Consequentialism. If we were also to accept Hedonism About Goodness, then we would be committed to supposing that it counts in favor of the conclusion that we ought to do a thing that our doing it will cause a man to feel pleased at the pain of another—and indeed, counts the more strongly in favor of this conclusion, the more pleased he will be. We may well think that must be wrong.

Well, I will be suggesting that we should reject Consequentialism, so the fact that if you accept it, you had really better not also accept Hedonism About Goodness does not strike me as a serious objection to Hedonism About Goodness. What is objectionable about Hedonism About Goodness is internal to it. Feeling pleased is feeling pleased *by* something, and there is a world of difference between pleasures according as their objects differ from each other.

So suppose we reject Hedonism About Goodness. If we wish to retain Consequentialism, we now have a problem on our hands.

4.

A serious problem. For which events are good and which bad? What answer to this question is to replace Hedonism About Goodness? Suppose my team plays your team in the football finals, and that my team wins. "That's good," I say. "That's bad," you say. Which of us is right? How on earth is that question to be answered?

It might be suggested that neither of us is right—that is, that the event of my team's winning is not itself either good or bad, that, as some philosophers would say, the event of my team's winning is not *intrinsically* good or bad. Rather, it is at most *instrumentally* good or bad, good

or bad only insofar as the events it will cause are intrinsically good or bad. Which, then, are the intrinsically good or bad events?

We might try to construct some examples. Suppose Alfred aims a gun at Bertha and fires it; Alfred misses, however, so Bertha survives. Perhaps we can say that Bertha's survival is an intrinsically good event. Bertha's death, had that occurred, would have been an intrinsically bad event; fortunately, her death did not occur.

There is a difficulty here, however, analogous to the one I pointed to when we looked at Hedonism About Goodness. For what if the reason why Alfred was aiming his gun at Bertha is that Bertha was villainously trying to kill Alfred? Let us suppose that Bertha, having survived, now kills Alfred, just as she had been villainously trying to. Are you still sure that Bertha's survival was an intrinsically good event? And that her death would have been an intrinsically bad one?

We might well want to say that a person's surviving—whether or not the person is a villain—is good for him, and that his death would be bad for him. But we need to remember that an event that is good or bad for one person can be the opposite, namely bad or good, for another person.

It has to be remembered also that what is in question here is not whether an event is good *for* or bad *for* a person, but rather whether it is just plain good or just plain bad. That is our question. And it is not in the least clear how it is to be answered.

5.

But what is good or bad for people must surely be in some way relevant to whether a person ought to act. Perhaps what is good or bad for people itself fixes what is just plain good or just plain bad? Perhaps a Consequentialist should therefore opt for the following idea about goodness: one event is better than another if and only if the first is 'more better for more' than the other. Consider again the event that consists in my team's winning its game with yours. "That's good," I say. "That's bad," you say. I asked: which of us is right? Perhaps the answer to this question is to be found out by finding out how many people the event is good for and how many it is bad for, and how good it is for those it is good for, and how bad it is for those it is bad for—the answer to the question being an appropriate function of those facts.

A Consequentialist who likes this idea could then say: the world will be better if a person does such and such than it will be if he does anything it is open to him to do instead just in case his doing the such and such will be more better for more. And if it will be, then it follows that that is what he ought to do.

There is a difficulty for this view that is a first cousin of one we have met twice before. Suppose it would profit me a lot to make you suffer a minor loss; suppose also that no one else would be affected by my act. It follows on this view that I ought to make you suffer the loss. That can't be right. There may well be cases in which it is permissible, even morally required, that one person cause another a loss. But it can't at all plausibly be thought that the mere fact that I would gain more by acting than you would lose counts in favor of its being the case that I ought to act.

The Consequentialist who rejects this idea about goodness can of course reject this outcome. He can remind us that he did not say that what matters morally is what is more good for more: what he said is that what matters morally is what is more good. And he can declare that some events that are more good for more may perfectly well be bad events. In particular, an event that consists in taking advantage of another for one's own profit may well be a bad event even if it is more better for more than any alternative open to the agent at the time.

It pays to stress this point. It seems nigh on a necessary truth that what a person ought to do is what would make the world be best—more good—than any alternative. That is why Consequentialism seems so attractive. It takes only a moment's reflection to see that it is not only not a necessary truth but false that what a person ought to do is what would make the world be more better for more than any alternative.

Of course a Consequentialist who rejects this idea about goodness owes us a better one. And what might that be?

6.

To summarize where we have come so far. I drew attention in section 4 to the attractiveness of Consequentialism, which is the idea that a person ought to do a thing if and only if the world will be better if he does it than if he does anything else it is open to him to do instead. And we were supposing that whether the world will be better turns on a com-

parison between the goodness or badness of the events that will occur if he does or does not choose the option. But which events are good and which bad? Once we have cut ourselves loose from Hedonism About Goodness, and from the idea I described in the preceding section, we are out at sea, adrift. It would be no surprise if people found themselves wondering how anyone could be supposed to have good reason for believing that a person ought to act in this way or that.

The point may be put another way. According to Consequentialism, the concept 'ought' reduces to the concept 'good.' If you want to know whether someone ought to do a thing, you need to ask what events will occur if he does it and what events will occur if he does anything else, and whether those events will be good or bad, and if so, how good or how bad. Let us now distinguish between two ways in which it can seem that we are at risk of having to become skeptics about morality. The first is this: it may be said that we just can't find out what all the events are that will occur if a person does or does not do a thing. Consider again Alfred's pressing the doorbell. I mentioned some of the events that that event will cause, but there are surely indefinitely many others that it will also cause. If we don't know which they all are, we can't even begin to assess the goodness or badness of all of them, and therefore can't find out whether Alfred ought to press the doorbell. I will call this shallow skepticism about morality. It is skepticism *about morality* because it is skepticism about the possibility of finding out what Consequentialism says must be found out if we are to find out whether judgments about what people ought to do are true. But it is shallow skepticism about morality because it is ultimately skepticism about matters of fact. The shallow skeptic says that if we could find out about the relevant matters of fact, then finding out about what people ought to do would be no problem.

Some Consequentialists have been shallow skeptics about morality, and contentedly so. I have G. E. Moore in mind in particular. According to Moore, we must just hope for the best: if we manage to do what we ought to do, that is just good luck for us. Other Consequentialists have not been contented at the prospect of having to become shallow skeptics about morality. No matter for our purposes.

For what we have in fact reached is the prospect of something markedly more worrisome, namely deep skepticism about morality. What we have reached is that even if we knew about all the events that will occur if a person acts and all that will occur if he does not, we are

still in epistemological trouble because we have found no satisfactory way of settling which of those events would be good and which bad. What looms is the fact-value gap, and it looks unbridgeable.

So what's to be done?

7.

I suggest that the reason why we find no satisfactory way of answering the question which events are good and which are bad is that there is no such question. Consequentialism requires that there be such a question, and that we be able to answer it if we are to be able to tell whether a person ought to do a thing. That, I suggest, is itself a conclusive objection to Consequentialism.

If someone draws our attention to a certain event—say, Alfred's pressing a certain doorbell—and asks us whether that was or would be a good event, or a bad event, or neither, we should not think “Ah, what a hard question”; we should instead ask ourselves whether we so much as understand what we have been asked.

Why do we think we do understand? Or anyway, why do so many moral philosophers think they do? One answer emerges clearly in G. E. Moore's *Principia Ethica*, with which twentieth-century Anglo-American moral philosophy began.

Moore said it is clear that some things are good, some are bad, and some are neither. Goodness, he said, is the property that all and only the good things have in common. That is the property that we would be ascribing to a thing—whether an event or anything else—if we said of it “That's good”; and that is the property such that we are asking whether a thing possesses it when we ask about the thing “Is it good?”

This idea seems to issue from nothing better than an oversimplified conception of the way in which the adjective “good” functions in English. When people say about a thing “That's good,” what they mean is always that the thing is *good in some way*. Perhaps they mean that the thing is a good fountain pen. Or a good book. Or a good apple. If so, what they mean is that the thing is good of a kind.

There is more too. A person might say “That's good,” not meaning that the thing is good of a kind, but that it is good for use in doing this or that. Perhaps that the thing is good for use in making cheesecake. Or they may mean that the thing is good *for* such and such or so and so. Per-

haps that the thing is or would be good for Alfred, or for England, or for the tree in my backyard. Or they may mean that the thing tastes good or looks good.

When talking about a person, they may say “He’s good,” meaning by this that he’s good at playing chess, or that he is morally good—just or honorable or generous. When talking about an experience or an activity, they may say “It’s good,” meaning by this that it’s pleasant or enjoyable.

What people *say* is the words “That’s good,” or “He’s good,” or “It’s good,” but what they mean—what *they*, but not their words, mean—is that the thing is good in one or other of the kinds of ways I have indicated.⁴ It is the context in which they assert those words that makes clear what they meant by the words, that is, what, perhaps given their preceding remarks, their hearers are entitled to suppose they mean. If the context does not make this clear, then their hearers are at a loss.

We should be clear that the ways in which a thing can be good that I have been indicating are not *grounds* for thinking a thing is good. St. Francis was good. How so? Well, he was a morally good person—he was just and kind. Chocolate is good. How so? Well, it tastes good. If what I have supplied you with are grounds for thinking that St. Francis and chocolate are good, that is, grounds for thinking that they both possess the property goodness, then it ought to be in order to ask which is better, for the adjective “good” has a comparative. But do you make sense of the question whether St. Francis was better than chocolate?

I think we had better conclude that there is no such property as goodness. All goodness, as we can put it, is goodness in a way. When it is asked whether a thing is good—whether the thing is a book or pie tin, or a person or an event—the context, or the speaker, needs to let us in on what the relevant way of being good is, or we not only can’t answer the question, we don’t even know what question was asked.

Consider events in particular. Suppose someone asks whether Alfred’s pressing the doorbell is or would be a good event. We should reply “How do you mean? Do you mean ‘Would it be good *for* somebody?’” And we had better be told whether that is what is meant, or whether something else is meant. We had better not be told that what is in question is instead whether the event is just plain, pure good, for there is no such thing.

⁴This is a point I have made in a number of other places, most recently in “The Right and the Good,” *Journal of Philosophy* 94, no. 6 (June 1997).

8.

Consequentialism, then, has to go. What is to replace it is a hard question, the harder in that Consequentialism rests on ideas that are very attractive. I will concentrate on it in Part II. Meanwhile, however, it pays to take note of some things we gain if we reject the question whether or not a thing is plain, pure good.

Most important, we are not now confronted with an unbridgeable fact-value gap.

For in the first place, there is no one fact-value gap: if there is one, there are many. Suppose we know a lot of facts about a certain fountain pen: how much ink it will hold, that it does not leak, how smooth its nib is, and so on. Ah, but is it a *good fountain pen*? Again, suppose we know a lot of facts about what a certain event would cause. Let it be the event that consists in Alfred's drinking some hot lemonade. Suppose, then, that we know, in particular, what that event would cause, given the condition Alfred is currently in—as it might be, that he has a sore throat. Ah, but would Alfred's drinking some hot lemonade be *good for him*? Again, suppose we know a lot of facts about how a certain brandy tastes: austere and delicate. (I take this description from a *New York Times* article on brandies.) Ah, but does the brandy *taste good*? If there are fact-value gaps, then I have drawn attention to three of them, for it is not at all plausible to think that what we have here are three cases in which we have facts in hand, and need to be told what consideration—the same in all three—would take us from the facts to the values. Whatever it is, if anything, that would entitle us to pass from those facts about the fountain pen to the conclusion that it is a good fountain pen is not at all plausibly thought to be the same as what would entitle us to pass from facts about Alfred and hot lemonade to the conclusion that his drinking some would be good for him, or the same as what would entitle us to pass from facts about how the brandy tastes to the conclusion that it tastes good.

Second, we should ask whether there really are unbridgeable fact-value gaps in the cases I mentioned.

What facts about a fountain pen warrant concluding that it is a good fountain pen? Well, some things are clear. It mustn't leak, it must be sturdy, it must hold enough ink to write several pages before filling, its nib must be smooth so as not to tear the paper being written on. A good fountain pen is one that would serve well the typical purposes of those

who want fountain pens. And whether a pen would serve those purposes well is something we can and do find out all the time.

If Alfred has a sore throat, then it is very likely that the event that consists in his drinking some hot lemonade would be good for him. Why so? Well, it is very likely that his drinking some will make him feel better. Of course, that might be mistaken. Perhaps he has an ulcer as well as a sore throat; then, perhaps, drinking hot lemonade would not make him feel better, and would in fact be bad for him. We know perfectly well what *kinds* of consideration bear on the question whether that event would be good for him. There certainly are cases in which it is hard to find out whether an event would be good for a person, and among them are cases in which we may have to conclude that we cannot find this out. Perhaps we are unable to attach weights to the various considerations that bear on whether the event would be good for him, as, for example, where the event consists in his making this or that choice among possible careers. Still, there are limits to what counts as a consideration in such cases.

Whether something tastes good is a messier matter. That is partly due to the fact that we have so little in the way of phenomenological characterizations of tastes—getting past “sweet,” “sour,” “bitter,” and “salty” is, for most of us, rather a stretch. It takes a professional to describe a brandy as austere and delicate. Moreover, most of us do not really attend to tastes very closely, and do not notice in them what a professional notices.⁵ When you think on how important the tastes of things are to us, that can seem very surprising. The questions that arise here are interesting and, I think, insufficiently studied by philosophers. For our purposes, however, it is perhaps enough to draw attention to the fact that there is a difference, which is plain to all of us, between a person’s *liking* the taste of something and its tasting good. Lots of people like the taste of strawberry Kool-Aid: it sells very well indeed. For all that, strawberry Kool-Aid does not taste good.⁶

There are ways of being good that are of particular interest to the

⁵ More from the *Times* article about another brandy: it is “round and notably spicy in flavor, with hints of nutmeg, cinnamon and hazelnuts.” J. L. Austin asked—in *Sense and Sensibilia*, ed. G. J. Warnock (Oxford: Oxford University Press, 1962)—“What sort of reception would I be likely to get from a professional tea-taster, if I were to say to him, ‘But there can’t be any difference between the flavors of these two brands of tea, for I regularly fail to distinguish between them?’” He left it to us to supply the answer.

⁶ There is a term that I think appropriate here. Lots of people like what can best be described as kitsch. Kool-Aid is kitsch in the realm of taste.

moral philosopher, and I will be returning to them in Part II. Meanwhile, it just is not in general true that our intellectual lives are everywhere crisscrossed by unbridgeable fact-value gaps. The adjective “good” is among the most commonly used in the English language. What we should have been doing is to look at how it is in fact used, and at what does in fact settle that it is or is not applicable.

9.

A further benefit can be got by attending to the ways of being good. There is a concept which has been much leaned on by many contemporary moral philosophers, but which has seemed very dark to others. What I refer to is the concept ‘reason for a person to do such and such.’ It has been thought to have an intimate connection with the concept ‘ought.’ Some philosophers hold that it is not the case that a person ought to do a thing unless there is a reason for him to do it. Or even more strongly: what a person ought to do is precisely what there is most reason for him to do. I leave aside for the time being the question how the concept ‘ought’ connects with the concept ‘reason for acting.’ Let us ask instead what must surely be the prior question: what *is* a reason for a person to do such and such?

It is easy enough to begin: a reason for a person to do a thing is something that counts in favor of his doing it. But what is that?⁷ There are a number of answers in the literature—I will discuss two of them.

Before we turn to them, however, we need to adopt two regimentations. The need for the first issues from the existence of a scatter of locutions in which the term “reason” appears: we need to decide how to connect the most common of them.

The weakest is the kind I started with, an example of which is:

- (1) There is a reason for Alfred to press the doorbell, namely X.

I take it that (1) is consistent with Alfred’s not believing that X is a rea-

⁷T. M. Scanlon says, “Any attempt to explain what it is to be a reason for something seems to me to lead back to the same idea: a consideration that counts in favor of it. ‘Counts in favor how?’ one might ask. ‘By providing a reason for it’ seems to be the only answer.”

And Scanlon therefore says he will take the concept of a reason as primitive. See his *What We Owe to Each Other* (Cambridge: Harvard University Press, 1998), 17. As will emerge in section 12 below, I think there is another, better because more informative, answer to the question what it is for X to count in favor of a person’s doing a thing than that X is a reason for the person to do it.

son for him to press the doorbell. Indeed, I take (1) to be consistent with Alfred's believing that there is no reason at all for him to press the doorbell.

There is room for dispute about what is the actual, or anyway the most common use of

(2) Alfred has a reason for pressing the doorbell, namely *X*.

For simplicity, I bypass arguments about usage. I will take it that Alfred can't have a reason for pressing the doorbell unless there really is one. Thus I will take it that (2) entails (1). On the other hand, I will also take it that it can't be the case that Alfred has a reason for pressing the doorbell, namely *X*, unless Alfred believes that *X* is a reason for him to press the doorbell. Thus I will take it that although (2) entails (1), (2) is not entailed by (1). I think that this decision does capture the most common use of (2), but whether it does does not matter for our purposes.

Suppose Alfred is now in process of pressing the doorbell. We may say:

(3) Alfred's reason for pressing the doorbell is *X*.

I take (3) to be stronger than (2). Alfred can have a reason for pressing the doorbell, namely *X*, and nevertheless not press it. (Perhaps he has a better reason for not pressing it.) By contrast, *X* can't be Alfred's reason for pressing the doorbell unless he is pressing it for that reason, and thus unless he is in fact pressing it. So (3) entails (2), but (3) is not entailed by (2).

So (3) is stronger than (2), and (2) is stronger than (1). I think that no theoretical issue turns on these decisions about (1), (2), and (3); that is why I say that what is in question here is (mere) regimentation.

Similarly for reasons for believing, or wanting, or expecting, or regretting, or hoping for, or feeling angry at, or . . . a thing; thus for anything for which there might be, and a person might have, a reason.

The need for a second regimentation is due to the fact that we need to fix on a general characterization of what a reason *is*. What, after all, might *X* be?

A reason is something one might reason *from*. What might a person reason from? Suppose that Alfred believes that Bertha's pig can fly. Why? Suppose the situation is this: Alfred believes that all pigs can fly, and therefore concludes that Bertha's pig can. We have three options.

(i) We can say that Alfred's reason for believing that Bertha's pig can

fly is the fact that he, Alfred, believes that all pigs can fly. While pigs can't fly, Alfred does anyway believe that they can.

I think it plain that this option is not a happy one. Alfred does not reason that Bertha's pig can fly from the fact that he, Alfred, believes that all pigs can. His reasoning, we are supposing, went like this: "All pigs can fly, therefore Bertha's pig can." His premise was not that he, Alfred, believes that all pigs can fly, but rather that all pigs can fly.

I am not suggesting that a person couldn't reason to a certain conclusion from the fact that he believes this or that. That does seem to be possible. Suppose that Charles has loved Dora for years, but his suit had always seemed hopeless. He is now suddenly struck by the thought that Dora loves him too. He concludes—from the very fact that he now believes she does—that there must have been some evidence of her love for him in her past behavior, evidence that was unrecognized by him at the time, and is still unclear to him now. Cases in which a person reasons to a conclusion from the fact that he believes this or that must surely be rare, however: normally, we reason not from our believing something, but rather from what we believe.

I did not spend time on option (i) because I think it a plausible description of Alfred: I did so because an analogue of the point I make here will reappear later.

A second possible description of Alfred, (ii), is that his reason for believing that Bertha's pig can fly is the proposition that all pigs can fly. That proposition is false, but we could say never mind: Alfred thinks it true, and reasons from it to his conclusion. More generally, we can say that a reason is always a proposition, true or false, which someone who thinks it true might reason from; and where a person thinks a proposition is true and reasons from it, *it is his* reason.

The third possible description of Alfred, (iii), is that Alfred has no reason for believing that Bertha's pig can fly. This is what we say if we take it that a reason is always a fact. Alfred might himself say "My reason for believing that Bertha's pig can fly is the fact that all pigs can fly." If that is what he says, then it is clear that although he thinks he has a reason for believing that Bertha's pig can fly, he doesn't actually have one, since there is no such fact as the fact that all pigs can fly.

It should be noticed that opting for (iii) is compatible with supposing that there is an answer to the question why Alfred believes that Bertha's pig can fly. We can say that he believes that Bertha's pig can fly because he believes that all pigs can. In other words, we can explain his

believing that Bertha's pig can fly. We cannot explain his believing that Bertha's pig can fly by giving his reason for believing it, since, according to (iii), he hasn't any; but he believes that all pigs can fly, and though that belief of his is false, the fact that he has it, we can say, is why he believes that Bertha's pig can fly.

I think that no deep theoretical issue turns on a decision between (ii) and (iii), for I think that any interesting claim we make about reasons on the supposition that they are propositions has an equally correct or incorrect analogue about reasons on the supposition that they are facts. That is why I take it, once again, that what is in question here is (mere) regimentation. Since taking reasons to be facts seems to me to square with usage better than taking them to be propositions, I will take them to be facts.

It should perhaps be stressed: I will be taking it that a reason is a fact not merely where it is a reason for believing something, but also where it is a reason for doing something. Indeed, also where it is a reason for feeling something or for wanting something and so on—that is, for whatever it is that a reason might be a reason for.

So much for regimentation. I said at the beginning of this section that a reason for a person to do a thing is something that counts in favor of his doing it. Given our second regimentation, we can re-put this point as follows: a reason for a person to do a thing is a fact that counts in favor of his doing it. Which facts are those? I said I would discuss two answers that may be found in the literature.

IO.

According to the first answer, every reason for action is a desire, or want.⁸ Suppose that Alfred wants to please Bertha. Suppose also that his pressing a certain doorbell would please her. Then, on this view, there is a reason for Alfred to press the doorbell, namely his wanting to please Bertha. What makes his wanting to please her a reason for him to press the doorbell is the fact that his pressing it would please her; but that being a fact, his wanting to please her *is* a reason for him to press it.

Given our first regimentation, it can of course be the case that although

⁸What I describe here is a simplified version of the theory argued for by Dennis Stampin in "The Authority of Desire," *Philosophical Review* (July 1987).

(1) There is a reason for Alfred to press the doorbell, namely his wanting to please Bertha,
is true,

(2) Alfred has a reason for pressing the doorbell, namely his wanting to please Bertha,

is false—after all, he may not know that his pressing the doorbell would please Bertha, and therefore not know that his wanting to please Bertha is a reason for pressing it. But if he does know that his pressing the doorbell would please her, and therefore believes that his wanting to please her is a reason for pressing the doorbell, then (2) is true. And we can understand how that reason counts in favor of his ringing the doorbell: it makes his ringing the doorbell be attractive to him.

This view is entirely compatible with our having bad reasons for action. Suppose instead that Alfred wants to annoy Bertha, and that his pressing the doorbell would in fact annoy her. Then there is a reason for him to ring the doorbell, namely his wanting to annoy Bertha. If he knows these things, then he has a reason for pressing the doorbell—and the reason counts in favor of his pressing the doorbell in that it makes his pressing the doorbell be attractive to him. That outcome is as it should be, for any theory of reasons for action must allow for the possibility of, indeed the fact of, there being bad reasons for action as well as good ones.

The theory also allows for there being stronger or weaker reasons for action: this difference (it says) turns on the strength or weakness of the want.

Given our second regimentation, however, this won't quite do as it stands. A desire is presumably a mental state. For example, and more precisely: Alfred's wanting to please Bertha consists in his being in a certain mental state. Alfred's wanting to annoy Bertha consists in his being in a different mental state. Is his being in a mental state a fact? Surely not: that idea seems to be a category mistake.

If a person's wanting something consists in the person's being in a certain mental state, then it looks as if no desire is a reason for action, since no desire is a fact.

An emendation all but suggests itself. A friend of this first theory could say that a reason for action is not itself a desire, but is instead the fact that consists in the person's having the desire.⁹ Thus the reason there is for Alfred to press the doorbell is the fact that Alfred wants to

⁹Stampe agrees that a reason is a fact, and he tells us that when he says that a desire is a

please (or annoy) Bertha. And it is that—that fact—that makes his pressing the doorbell be attractive to him.

But the theory still won't do. As I said earlier, a reason is something a person might reason *from*. There surely are cases in which a person reasons to something from the fact that he wants something. In particular, there are cases in which a person takes the fact that he wants something to be a reason for doing something. Suppose Carol has always disliked milk in the past. She now finds herself wanting to drink some. There is a theory according to which people's wants in food often issue from a nutritional deficiency, thus, for example, that wanting to drink some milk might issue from a need for calcium. Let's suppose that theory is true, and that Carol believes it is. Then she might think: "Here is an interesting fact I have just noticed about myself: I want to drink some milk. Wanting to drink milk sometimes issues from a need for calcium. Therefore I may well need some calcium (which milk supplies). So my wanting to drink some milk is a reason for me to drink some." This is probably a relatively rare kind of case, however: normally, we reason not from the fact that we want something, but rather from facts about what we want and how it might be got.

Carol, who takes the fact that she wants to do something to be a reason for doing it, should remind us of Charles, who takes the fact that he believes one thing to be a reason for believing another. Both kinds of case are surely rare.

If so, then we cannot suppose that every reason for action is a desire. At most some are.

I I .

I said that I would discuss two answers to the question what facts are reasons for action. Fortunately we can be brief about the second.

On this second view, every reason for action is a combination (a pair? a conjunction?) of a desire and a belief.¹⁰ Suppose that Alfred wants to please Bertha and believes that his pressing a certain doorbell would

reason for action, what he really means is this: the fact that the person wants something is a reason for the person to act accordingly. See *ibid*.

¹⁰An influential example is Donald Davidson's account of reasons for action in "Actions, Reasons, and Causes," *Journal of Philosophy* 60 (1963). That article was reprinted in his *Essays on Actions and Events* (Oxford: Oxford University Press, 1980). According to Davidson, any pro-attitude can play the role of the required desire. Friends of this theory, and it has a great many, do not typically trouble themselves over the question what the mode of combination is, and I leave it open.

please her. Then, on this view, there is a reason for Alfred to press the doorbell, namely the combination consisting of his wanting to please Bertha and his belief that his pressing it would please her. Moreover, Alfred has a reason for pressing the doorbell, namely that want/belief combination—which counts in favor of his pressing the doorbell since it makes doing so be attractive to him.

We can, as I said, be brief about this theory. Given our second regimentation, a reason is a fact, so we had better reconstrue the theory to say that a reason for acting is not a want/belief combination but rather a combination consisting of the fact that a person wants this and the fact that he believes that. A reason, however, is something that one might reason from. But as we know, it is rare for a person to reason to something from the fact of his wanting this, and rare for a person to reason to something from the fact of his believing that—remember the cases of Carol and Charles—and presumably therefore at least as rare for a person to reason to something from the fact of his wanting this *and* believing that.

Why are so many philosophers inclined to say that a reason for acting is a want/belief combination? I think it pretty clear why they do. First, they think that a person's reason for doing a thing explains his doing of it.¹¹ Why so? It seems very plausible to think that if Alfred is pressing a doorbell for a reason, then when we have found out what his reason was, we have found out why he pressed it. That is, we have found an explanation of his pressing it. Indeed, what we have found, namely his reason for pressing the doorbell, itself explains his pressing it.

Following Hume, as many contemporary philosophers do, we may well believe, second, that something explains a person's doing a thing only if it contains a want and a belief. Something explains Alfred's pressing the doorbell only if it contains a want and a belief, in particular, the want and belief because of which he pressed it.¹²

¹¹That this motivates the theory emerges clearly in Davidson, "Action, Reasons, and Causes." I hazard a guess that it also motivates the theory offered in Stampe, "The Authority of Desire."

¹²Hume did not, and contemporary philosophers also do not, say that something explains just anyone's doing just anything only if it contains a want and a belief: on their view, this holds only of doings that are intentional. (If a man is nervously and unwittingly tapping his fingers, then it may be that there is no want and belief out of which he is doing it; and his doing of it is presumably explainable by appeal to something else, perhaps in a case such as this, to anxiety.) On the other hand, a person who does a thing for a reason does do it intentionally, and it is only doings of things for reasons that concern us here. I therefore omit the qualification.

It follows that a person's reason for acting contains a want and a belief: a reason for acting is a want/belief combination.

This won't do, however. Following Hume, what explains a person's doing a thing is his wanting this and believing that. But the fact that he wants this and believes that is not likely to be his reason for doing the thing: such cases are rare. The more usual case is like this. We ask Alfred why he is pressing the doorbell, and he replies: "to please Bertha." We find out from his saying this that he wants to please Bertha, and believes that his pressing the doorbell will please her. And that—the compound fact that he wants to please Bertha and believes that his pressing the doorbell will please her—does explain his pressing the doorbell. But that compound fact is not his reason for pressing it.

This is why we can find out why a person does a thing even if he has no reason for doing it, but only thinks he does. Consider, first, reasons for belief. I invited you to imagine that Alfred believes that Bertha's pig can fly. We ask him why he believes this, and he says "All pigs can fly." According to our second regimentation, Alfred has no reason for believing that Bertha's pig can fly since there is no such fact as the fact that all pigs can fly. But we can nevertheless explain his believing that Bertha's pig can fly: he believes this because he believes that all pigs can.

So similarly for reasons for action. Suppose we ask Alfred why he is pressing a certain doorbell, and he replies: "to please Bertha." But suppose that his pressing the doorbell will not in fact please Bertha; suppose it will instead annoy her. Then Alfred has no reason for pressing the doorbell. But we can nevertheless explain his pressing it: he is doing so because he wants to please Bertha, and believes that his doing so will please her.

12.

So what is a reason for doing a thing? It is something that counts in favor of doing it. Given our second regimentation, we can re-put this point as follows: a reason for a person to do a thing is a fact that counts in favor of his doing it. Which facts are those? We have now rejected two answers that may be found in the literature.

I suggest that it would pay us to help ourselves to the fact that all goodness is goodness in a way, and to consider the following sufficient condition for a fact to be a reason for a person to do a thing: a fact is a rea-

son for a person to do a thing if it is a fact to the effect that his doing the thing would be good in some way.¹³ There is a reason for Alfred to press that doorbell if his pressing it would be good for him. Or good for Bertha or Charles. Or would be enjoyable. (I suppose a person might like pressing doorbells.) And so on. And I suggest that it is clear that a fact to the effect that Alfred's pressing the doorbell would be good in some way counts in favor of his pressing it.

The fact that Alfred's pressing the doorbell would be good in some way may not everywhere count very strongly in favor of his pressing it: its being enjoyable to ring doorbells does not count very strongly in favor of pressing yours or mine. But the fact that doing so would be enjoyable does count in favor of doing so.

Moreover, Alfred has a reason for pressing the doorbell if, for some way of being good, W , he knows that his pressing the doorbell would be good in way W . His knowing that pressing it would be good in way W is, of course, entirely compatible with his not pressing it. For example, he might not care in the least about the fact that his pressing the doorbell would be good in way W . But suppose he does care: suppose, in particular, that he wants to do something that is good in way W , and, in fact that he wants to do something that is good in way W more than he now wants to do anything else. Then given that he also believes that his pressing the doorbell would be good in way W , he may be expected to press it.¹⁴ If he does press it, then his reason for pressing it is the fact that his doing so would be good in way W ; and his wanting to do something good in way W , together with his believing that pressing the doorbell is good in way W , explains his pressing the doorbell.

It is perhaps worth stressing that Alfred can believe that his pressing the doorbell would be good in a way W compatibly with (i) his not wanting to press it, and (ii) his not wanting to do something good in way W . That his believing that his pressing the doorbell would be good in way W is compatible with (i), his not wanting to press it, is probably obvious enough. Even if his belief is that his pressing the doorbell would be good for him. I may know that my taking a certain nasty-

¹³G. E. M. Anscombe said that a bit of practical reasoning must be understood to have a major premise that attributes what she called a "desirability characterization" to the act contemplated. See her *Intention* (Oxford: Basil Blackwell, 1957). I suggest that we can interpret her as meaning that the major premise must assert that the act would be good in some way.

¹⁴I do not say that he will press it: I say only that he may be expected to do so, since I wish to leave room for the possibility of weakness of will.

tasting medicine would be good for me without wanting to take it. If I do take it, then that is presumably because, given that taking it would be good for me, I am willing to take it; but being willing to take it is not the same as wanting to take it. I hope it is also obvious that Alfred's believing that his pressing the doorbell would be good in way *W* is compatible with (ii) his not wanting to do something good in way *W*. Even if the way of being good is goodness for him. I may refrain from taking the nasty-tasting medicine, not merely because it tastes bad, but because I do not now want to do what is good for me.

The fact that one can know that doing a thing would be good in a certain way compatibly with wanting neither to do the thing nor to do something good in that way is due to the fact that it is *goodness in a way* that we are dealing with. A number of philosophers have held that believing a thing would be pure good, intrinsically good, motivates the believer to try to bring it about.¹⁵ In particular, then, they have held that believing that one's doing a thing would be good motivates one to do it. Other philosophers have disagreed. Given the illegitimate abstraction that issued in the idea that there is such a property as goodness, and the consequent unclarity about what the property is, it is perhaps no surprise that this disagreement resists resolution.

In sum, I suggest that it would pay us to consider the following sufficient condition for a fact to be a reason for a person to do a thing: a fact is a reason for a person to do a thing if it is a fact to the effect that his doing the thing would be good in some way. Indeed, I think it a very plausible idea.

13.

There is a possible objection: accepting that sufficient condition yields that for each of the many things you might now do, there probably is at least one reason for you to do it. That is because there are many ways of being good, and for each of the many things you might now do, there

¹⁵ See, for example, Charles L. Stevenson, "The Emotive Meaning of Ethical Terms," *Mind* 46 (1937), reprinted in his *Facts and Values* (New Haven: Yale University Press, 1963). J. L. Mackie's reason for suspicion of the property goodness is the very fact—he takes it to be a fact—that if there were such a property, it would have to be a property such that believing that a thing would have it motivates the believer to try to bring the thing about. See his *Ethics: Inventing Right and Wrong* (London: Penguin Books, 1977).

probably is at least one way in which your doing it would be good. Your mowing my lawn would be good for me. Your watering my lawn would be good for me and it would also be good for my lawn. Suppose your neighbor is doing something illegal; your helping him to hide the traces might be good for him. Again, your going out to dance a jig in the street might be good for use in a list (currently being compiled) of eccentric behaviors in the suburbs. And so on. But (so the objection goes): that many reasons for doing that many things is too many.¹⁶

My own impression is that this is not a worrisome objection. A reason for doing a thing counts in favor of doing it; but there being a fact that counts in favor of your doing a thing is entirely compatible with its being of no interest to you that there is. There being such a fact is also compatible with its being wrong for you to do the thing.

Moreover, any of the facts to the effect that your doing a thing would be good in a way might be your reason for doing it. As I say, you might not care that your doing a thing would be good in way *W*. But you might care; and it might be that you therefore do the thing for that reason.

It could of course be insisted that a fact to the effect that your doing a thing would be good in way *W* is a reason for you to do it only if you want to do what would be good in way *W*. Thus suppose you don't care that your watering my lawn would be good for me; suppose what you care about is lawns—what you want is that *they* be in good condition. Then it might be said that the fact that your watering my lawn would be good for me is not a reason for you to water my lawn; rather, it is only the fact that your watering my lawn would be good for the lawn that is a reason for you to water it.

If we accept this idea, then we have to say that the fact that your doing a thing would be good in way *W* is a reason for you to do it only if you want to do what would be good in way *W*. Moreover, that such a fact might be a reason for you to do a thing at one time and not another. That would be the case if at the earlier time you did not want to do what would be good in way *W* and then came to want to do so at a later time. Perhaps some people would find this narrow construal of reasons for action attractive.

I see no theoretically important reason for rejecting this narrow construal. Or for accepting it. Whatever we wish to accomplish in moral theory can be—indeed, had better be—accomplishable whether we opt

¹⁶I thank Sarah Stroud for making this objection: she said that on this view, reasons for action “come too cheap.”

for this narrow construal of reasons for action or the broad construal of them that I have recommended.¹⁷ I will return to this matter in Part II.

Meanwhile, I think it clear that the broad construal that I have recommended is more squarely in accord with our ideas about reasons for action. Consider a man who is standing by, watching a child drown. He has a life preserver, and could easily throw it to the child. His doing so would be good for the child. On the other hand, he cares not the least about the child, and does not want to do what would be good for it. I am inclined to think it simply false to say that there is no reason for him to throw the life preserver. There is a reason for him to do this, lying in the fact that his doing it would be good for the child. His not wanting to do what would be good for the child does not mean that there is no reason for him to do it; it means merely that he is a thoroughly bad person.

In sum, I suggest that we should reject this objection. We should agree that a fact is a reason for a person to do a thing if it is a fact to the effect that his doing it would be good in a way. Whether the person will do a thing there is reason for him to do, or ought to or must do it, are quite other matters, fixed by quite other considerations.

14.

Is the condition I have offered necessary as well as sufficient? One option is to say that it is, thus to say:

- (i) A fact is a reason for a person to do a thing if and only if it is a fact to the effect that his doing it would be good in a way.

Should we agree? No. It is surely plain that the fact that your *not* doing a thing would be bad in a way is a reason for you to do the thing. (The fact that your not doing the thing would be bad in a way certainly counts in favor of your doing it.) But the fact that your not doing the thing would be bad in a way is not itself a fact to the effect that your doing it would be good in a way.

Second, we should also allow that the fact that your doing a thing

¹⁷The narrow construal is recommended by Bernard Williams; see his "Internal and External Reasons," which first appeared in *Rational Action*, ed. Ross Harrison (Cambridge: Cambridge University Press, 1980), and was reprinted in Bernard Williams, *Moral Luck* (Cambridge: Cambridge University Press, 1981). See also T. M. Scanlon's careful discussion of Williams's arguments in Scanlon *What We Owe to Each Other*, App. Scanlon prefers the broad construal, as do I, but I take him to agree that there is no theoretically important reason for preferring it to the narrow one.

would be better in a way than your doing anything else is a reason for you to do the thing. (This fact too counts in favor of your doing the thing.) But this fact is not itself a fact to the effect that your doing the thing would be good in a way.

Third, we should also allow that the fact that someone has a right that you do the thing is a reason for you to do it. But this fact too is not in itself a fact to the effect that your doing the thing would be good in a way.

These facts are all evaluative, or normative, and no doubt there are other examples that could be added. I suggest that what we should say is this:

- (ii) A fact is a reason for a person to do a thing if and only if it is a fact to the effect that
- his doing it would be good in a way, or
 - his not doing it would be bad in a way, or
 - his doing it would be better in a way than his doing anything else, or
 - someone has a right that he do it, or . . .

leaving room for other evaluative facts to be added.

But only other evaluative facts. For if we suppose that a reason for a person to do a thing is *itself* something that counts in favor of his doing it, then we should limit reasons for action to evaluative facts, as (ii) is intended to do. It is their being evaluative that marks these facts as reasons for action—since it is in virtue of their being evaluative that they count in favor of an action.

But perhaps we do not use the term “reason for action” as strictly as option (ii) requires? Suppose that Cora reasons to herself as follows: “My drinking some milk would increase my calcium intake, so I’ll drink some.”¹⁸ And suppose that her premise is true, thus that her drinking some milk really would increase her calcium intake. Let us call that the Calcium Fact:

(The Calcium Fact) Cora’s drinking some milk would increase her calcium intake.

She takes the Calcium Fact to be a reason for her to drink some milk; isn’t it plausible to think that it is? If it is, then—since the Calcium Fact is not evaluative—there are facts that are reasons for action that are not evaluative, and option (ii) must be rejected.

¹⁸Notice how Cora differs from Carol of section 10, whose reasoning rested on the premise that she wanted to drink some milk.

On the other hand, there is an assumption that Cora is making, namely, that her increasing her calcium intake would be good for her. (Else she would not take the Calcium Fact to be a reason for drinking some milk.) And if that assumption is false, then the Calcium Fact is not a reason for her to drink some milk. Thus, if her increasing her calcium intake would not be good for her, or would be positively bad for her, then the Calcium Fact is no reason at all for her to drink some milk.

We could say that the Calcium Fact is a reason for her to drink some milk if, but only if, her increasing her calcium intake would be good for her. If we do, we must reject the relatively simple option (ii). Alternatively, we could say that the Calcium Fact is not *itself* a reason for her to drink some milk. (It does not *itself* count in favor of her drinking some milk.) No doubt the Calcium Fact is a reason for believing that her drinking some milk would be good for her; after all, most of us really would benefit from increasing our calcium intake. We could therefore conclude that the Calcium Fact is (merely) a reason for believing that there is a reason for her to drink some milk. And we would therefore be able to retain the relatively simple option (ii).

My own impression is that nothing theoretically important turns on which of these options we choose, and I therefore recommend that we choose the relatively simple option (ii).

15.

My main concern in sections 9 through 14 has been the concept ‘reason for action.’ I have wished to bring out that we have an answer to the question “What is a reason for a person to do such and such?” if we help ourselves to the fact that all goodness is goodness in a way, and attend to the ways of being good.

I want to say a few things about desires by way of conclusion to Part I. I think it pays us to do so because here too it pays us to attend to the ways of being good.

We typically have a reason for wanting something, and I suggest that we should take reasons for wanting something to be similar to reasons for doing something. Thus I recommend that we say first: a fact *F* is a reason for a person to want a state of affairs *S* to obtain just in case *F* counts in favor of *S*’s obtaining. And that we say second: *F* counts in favor of *S*’s obtaining just in case *F* is a fact to the effect that *S*’s obtaining

would be good in a way, or *S*'s not obtaining would be bad in a way, or . . . , where the continuation is analogous to that in (ii) of section 14.¹⁹ If there is no such fact, then though the wanter may think there is a reason for him to want *S* to obtain, he is mistaken: there isn't.

So a person can want something and there be no reason for him to want it. Can a person want something without even believing there is a reason for him to want it? G. E. M. Anscombe invites us to imagine a man who tells us he wants a saucer of mud. "How so?" we ask. "What would be good about your getting a saucer of mud?" "Nothing," he replies; "I just happen to want to get one." Anscombe suggests that this is unintelligible, and I think she is right.²⁰ I think that you can't want something without thinking there is a reason for you to want it—just as (I should think) you can't expect or regret something without thinking there is a reason for you to expect or regret it.

If these ideas are right, then—even apart from the considerations I drew attention to in section 10 above—it would be no surprise if some philosophers thought that the fact that a person wants something is a reason for him to try to get it. For suppose that you want a state of affairs *S* to obtain. Suppose that you have a reason for that want, namely the fact that *S*'s obtaining would be good in a certain way *W*. A little piece of reasoning takes you from that fact to the conclusion that your trying to get *S* to obtain would also be good in way *W*; and *that* is a reason for you to try to get *S* to obtain. So your reason for wanting lends weight to a reason for acting, and it is therefore easy to think that the wanting is itself a reason for acting. Easy, but I suggest mistaken.

16.

I have argued in Part I that Consequentialism must be rejected on the ground that it reduces what a person ought to do to the maximizing of goodness, whereas there is no such thing as goodness. All goodness, I said, is goodness in a way. I then said that we gain more from attending to the fact that all goodness is goodness in a way than merely a refuta-

¹⁹Note the availability of a more complicated condition on reasons for wanting a thing, analogous to the one available for reasons for acting; that is, one that allows for the possibility that nonevaluative facts may be reasons for wanting a thing.

²⁰Anscombe, *Intention*. Warren Quinn makes a similar point in "Putting Rationality in Its Place," *Morality and Action* (Cambridge: Cambridge University Press, 1993). See also Richard Kraut, "Desire and the Human Good," *Proceedings and Addresses of the American Philosophical Association* 68; no. 2.

tion of Consequentialism: we are able to give an account of what it is for something to be a reason for acting.

What we should turn to now is what a person ought to do. I will not even try to produce a theory about what a person ought to do. I will only make some suggestions about the structure that I think such a theory should have.

PART TWO: ADVICE

I.

The word “ought” is probably just about as commonly used as the words “good” and “bad” are. When we say such things as “Alfred ought to drink some hot lemonade” or “Alfred ought to pay Bertha five dollars,” what does or would make what we say true? These assertions have a common form, which I will write

A ought to V

—they are obtainable from that expression by replacing somebody’s name for “A” and some verb or verb phrase for “V.” (I should point out that what replaces “V” may be the likes of “refrain from paying Bertha” as well as the likes of “pay Bertha.”) So what we are asking is: what does or would make an assertion of that form true?

It is certainly plausible to think that what a person ought to do is intimately connected with what would be good or bad. But assuming that I was right to say, as I did in Part I, that there are no such things as goodness and badness, or betterness or worseness, we cannot say

A ought to V if and only if A’s V-ing would be good

or

A ought to V if and only if A’s V-ing would be better than A’s doing any of the other things it is open to him to do.

A’s V-ing can’t be *just* good: it is at best good in this or that way or ways. And *A’s V-ing* can’t be *just* better than *A’s X-ing*, or *A’s Y-ing*, or *A’s Z-ing*: it can at best be better than those in this or that way or ways.

These considerations might tempt one to agree with those who say that the word “ought” is at least three ways ambiguous.

On any view, the word “ought” is at least two ways ambiguous. We say “The train ought to arrive by 3:00,” and when we do, we are not saying about the train what we say about Alfred when we say he ought to drink some hot lemonade. What we say about the train is roughly that the train may be expected to arrive by 3:00; what we say about Alfred is roughly that it is advisable that he drink some hot lemonade. I will throughout be using the word “ought” in the latter sense, which might be called the advice sense.

Now some people say that there is no such thing as *the* advice sense of the word “ought”: they say that “ought” has at least two advice senses. Suppose I ask whether Alfred ought to pay Bertha five dollars. They would reply: “What do you mean? Do you mean ‘would it be good for Alfred to pay Bertha five dollars?’ If that’s what you mean, then maybe the answer is no. Or do you mean ‘would it be morally good for Alfred to pay Bertha five dollars?’ If that’s what you mean, then maybe the answer is yes.” And if I say “Look, what I asked was just, simply, whether Alfred ought to pay Bertha five dollars,” they reply that there is no such question—no such thing as its *just* being or not being the case that a person ought to do a thing.

This idea might seem to square well with what I have been saying about goodness and badness. Thus it might be said that “ought” is multiply ambiguous: that it has a different meaning for each way in which an act can be good. In particular, it might be said that what we need is the following:

A ought_W to V if and only if A’s V-ing would be better in way W than A’s doing any of the other things it is open to him to do instead.

So, for example, that we should say about Alfred:

Alfred ought_{good for Alfred} to pay Bertha if and only if Alfred’s paying Bertha would be better for Alfred than Alfred’s doing any of the other things it is open to him to do instead,

Alfred ought_{morally good} to pay Bertha if and only if Alfred’s paying Bertha would be morally better than Alfred’s doing any of the other things it is open to him to do instead,

and no doubt also

Alfred ought_{enjoyable} to pay Bertha if and only if Alfred’s paying Bertha would be more enjoyable than Alfred’s doing any of the other things it is open to him to do instead,

and so on for each of the ways of being good. A proponent of this idea says that if anyone asks whether Alfred ought to pay Bertha, we cannot answer the question; indeed, no question has even been asked, unless what the speaker means to be asking is, for some particular way of being good W , whether Alfred ought _{W} to pay Bertha. We can call this the Multiple Ambiguity Idea.

It is in one way an attractive idea: as I said, it is plausible to think that what a person ought to do is intimately connected with what would be good or bad, and this idea expresses that connection in a very simple way—while accommodating the fact that all goodness is goodness in a way.

On the other hand, it is a very bad idea, and that for two reasons. First, it is wildly implausible to think that the word “ought” is multiply ambiguous in that way. Second, it gives moral goodness and badness the wrong role in assessments of what a person ought to do. I will return to moral goodness and badness later. Let us first look at the proposed ambiguity.

Suppose that Alfred is ill, and that only a dose of a certain medicine will cure him. It tastes truly awful, however. Alfred asks us “Ought I really take it?” It is a wildly implausible idea that we can reply only: “Well, your taking it would be very unpleasant, so in one sense of ‘ought,’ it’s not the case that you ought to take it, namely the ‘ought_{enjoyable}’ sense of ‘ought.’ But your taking it would be good for you, so in another sense of ‘ought,’ you ought to take it, namely the ‘ought_{goodness for Alfred}’ sense of ‘ought.’” It is likely that Alfred will repeat his question: “But ought I take it?” It surely won’t do to reply: “Are you deaf? I just told you that in one sense you ought to and in another sense it is not the case that you ought to, and that’s all the advice that anyone can give you.” We *can* give more advice: we can say what the case presumably warrants saying, namely that he ought to take the medicine.

Similarly for cases in which what is good for one is bad for another. Suppose that Alfred’s paying Bertha five dollars would be good for Bertha but bad for Alfred. Alfred asks whether he ought to pay Bertha five dollars. A proponent of the Multiple Ambiguity Idea says that the only advice we are in a position to give him is that in the ‘ought_{goodness for Bertha}’ sense of “ought” Alfred ought to pay Bertha, whereas in the ‘ought_{goodness for Alfred}’ sense of “ought,” it is not the case that Alfred ought to pay Bertha. Surely there remains a further question, namely whether Alfred just plain ought to pay Bertha.

In asking the question whether Alfred ought to take the medicine or ought to pay Bertha, we are using the word “ought” in *the* advice sense. That sense is what is sometimes called the ‘all things considered’ sense of the word. I think that “ought” has no non-all-things-considered sense.¹ But I will not argue for that here; I will merely use it in its all-things-considered sense throughout what follows.

I am going to make four suggestions that bear on it. As I said at the end of Part I, I will not offer a theory about what people ought to do. What I will do is just to draw a sketch of the structure of the theory I think we need to replace Consequentialism with. Alas, a theory with this structure would lack the simplicity of Consequentialism. But then we really shouldn’t have expected a theory of what we ought to do to be simple as Consequentialists take it to be.

2.

I said it is plausible to think that what a person ought to do is intimately connected with what would be good or bad. It is very plausible to think that it is certain particular ways of being good and bad that we should be attending to.

Let us begin with the easiest cases, namely those in which what I will call (Isolation)—

(Isolation) *A*’s *V*-ing would neither be bad for anything other than *A*, nor infringe anything’s rights—

is true. They include cases in which *A*’s *V*-ing would be bad in one or more ways. Indeed, they include cases in which *A*’s *V*-ing would be bad for him. They exclude cases in which, while *A*’s *V*-ing would be bad for nothing other than *A*, it would infringe a right of *B*’s. (This is surely possible: it might be the case that *B* has a right that *A* not *V*, though it would not be bad for *B* or anything else if *A* *V*-ed.)

The first of my four suggestions, then, is the following:

(I) If (Isolation) is true, and if also *A*’s *V*-ing would be better for him than his doing any of the other things it is open to him to do instead, then it follows that *A* ought to *V*.

I stress: I am suggesting that it *follows* that *A* ought to *V*. In particular,

¹ Some philosophers think that “ought” does have a non-all-things-considered sense; I describe their view very briefly in section 8 below—see note 10.

there is no need to add a further premise to the effect that *A* wants to do what it would be best for him to do. A person's wants do have a bearing on what he ought to do, but their bearing is indirect; we will return to the bearing of wants in the following section.

Another way to put that point is this. On some views, we should distinguish between categorical and hypothetical 'imperatives.' I put scare-quotes around the word "imperatives" since friends of those views would have us distinguish not merely between imperatives, such as "Drink that medicine" and "Pay Bertha five dollars," but also between sentences of the form that interests us, namely the likes of "Alfred ought to drink that medicine" and "Alfred ought to pay Bertha five dollars." On one interpretation of those views, then, some sentences of our form are categorical: one who asserts them asserts that the person named is to do the thing (drink the medicine, pay Bertha five dollars) whatever his wants may be. Others are hypothetical: one who asserts them asserts only that if the person wants most strongly what he would get in or by doing the thing, then he is to do it.² Then the point I made in the preceding paragraph can now be re-put as follows: what follows from the fact (supposing it a fact) that Alfred's drinking the medicine would be best for him, and bad for nothing other than Alfred, and would infringe no one's rights, is not a hypothetical but a categorical imperative—not that he ought too drink it if he wants most strongly to do what's best for him, but, simply, that he ought to drink it.

If Alfred does not at all want to do what is best for him, then it may well be, of course, that he will not drink the medicine. So also if, while he wants to do what is best for him, he wants more strongly to avoid the medicine's nasty taste. Either way, more's the pity for him—he is behaving imprudently.

Several considerations need to be mentioned before we move on. My first suggestion was that if (Isolation) is true, and also *A*'s *V*-ing would

²This is only one interpretation of the distinction between categorical and hypothetical 'imperatives.' On the interpretation I supply in the text above, the distinction is between two kinds of sentences. On another interpretation, the distinction is between two kinds of propositions a person might be asserting in asserting any given sentence of our form. Presumably there are other ways of interpreting the distinction. For our purposes, these differences do not matter.

My own view on this matter is that no such distinction can plausibly be made: that is, there is no use of the English word "ought" according to which one who asserts "A ought to *V*" asserts what is called a hypothetical imperative on any interpretation of that term. If I say "You ought to put antifreeze in your car in the winter," and it then turns out that you don't want to protect your car, and indeed, that it would be bad for you to protect your car, I don't say "Still, you ought to put antifreeze in your car in the winter, though of course I mean only to be asserting a hypothetical imperative." What I say is rather: "Sorry, I was mistaken in thinking you ought to put antifreeze in, etc. etc."

be better for him than his doing any of the other things it is open to him to do instead, then it follows that *A* ought to *V*. What if there is a tie?—that is, what if *A*'s *V*-ing and *A*'s *X*-ing would be equally good for him, and both better for him than his doing any of the other things it is open to him to do instead? I think that we should say here what, in Part I, I said the Consequentialist would say by way of reply to the same question. Thus we should say that in case of such a tie, it is not the case that *A* ought to *V* and not the case that *A* ought to *X*; rather, he may do either, and ought not do anything other than either.

A second consideration that needs to be mentioned is this: my first suggestion does not imply only that whether *A* wants to do what it would be best for him to do is irrelevant to the question whether he ought to *V*: it implies that *A*'s knowing, or even believing, that his *V*-ing would be best for him is also irrelevant to the question whether he ought to *V*. But that, I think, is exactly as it should be. Suppose, for example, that Alfred will be cured if and only if he drinks a certain medicine, and that his doing so would be bad for nothing other than himself, and infringe no one's rights. But suppose that while we know this, Alfred does not—suppose he not only does not know it but does not even believe it. Is it true, all the same, that he ought to drink the medicine? I recommend that we say that it is true. For suppose he asks us whether he ought to drink it. It would be weird to reply that we can't tell whether he ought to drink the medicine until he first tells us what he believes about it and about himself. We don't need to find out what he believes: whatever his beliefs may be, he ought to drink it. He has asked us for advice, and that is the advice we should give him.

Similarly, my first suggestion implies that what *A* would be intending in *V*-ing (if he *V*-ed) is also irrelevant to the question whether he ought to *V*. That too is as it should be. If Alfred asks us whether he ought to drink the medicine, it would be equally weird to reply that we can't tell until he first tells us what he would be intending in drinking it. Whatever he intends, he ought to drink that medicine.

I think it clear quite generally that what fixes whether a person ought to do a thing is not the person's subjective state of mind, but instead the objective facts of the situation he is in. I will summarize this point by saying that the question what a person ought to do is (not subjective, but rather) objective.³

³For more on this point, see Jonathan Bennett's distinction between first-order and second-order morality in his *The Act Itself* (Oxford: Oxford University Press, 1995).

There is a third consideration that needs to be mentioned. I said that we were to be considering cases in which

(Isolation) *A*'s *V*-ing would neither be bad for anything other than *A*, nor infringe anything's rights

is true. Why "not be bad for anything other than *A*" rather than "not be bad for any person other than *A*"? I meant to be setting aside cases in which *A*'s *V*-ing would be bad for anything. An animal, for example. Or for a living thing of some other kind, such as a plant. Or for some non-living thing, such as an artifact or a river. There obviously is such a thing as an act's being bad for an animal or a plant; if less obvious, it is also true that an act may be bad for an artifact or a river. Wherever *A*'s *V*-ing would be bad for something other than *A*, then that is in one or another way relevant to whether he ought, or even may, *V*: the question whether *A* ought to *V* is not settled by the mere fact that it would be best for him to do the thing and bad for no other person.

Again, why "not infringe anything's rights" rather than "not infringe any person's rights"? I meant to leave open that things other than people have rights. Organizations. Animals, perhaps. Wherever *A*'s *V*-ing would infringe anything's rights, then that is in one or another way relevant to whether he ought, or even may, *V*: the question whether *A* ought to *V* is not settled by the mere fact that it would be best for him to do the thing and infringe no person's rights.

We will turn to cases in which (Isolation) is not true in section 5.

3.

It is intuitively clear that the concepts 'good for' and 'bad for' play an important role in fixing whether a person ought to do a thing. We should therefore have a look at them.

There is a large literature on what it is for something to be good for a person, but I think that any plausible theory will serve our purposes well enough.

I assume three constraints on the plausibility of a theory about what it is for something to be good for a person. First, the theory must allow for the possibility that a person's doing a thing would be good for him, even best for him of all his alternatives, and yet bad for another person. Similarly, it must allow for the possibility that a person's doing a thing

would be bad for him, even worst of all his alternatives, and yet good for another person. That this is a constraint on the plausibility of a theory is probably obvious enough.

A second constraint is stronger: the theory must allow for the possibility that a person's doing a thing would be good for him, even best for him of all his alternatives, though he ought not do it; and that his doing a thing would be bad for him, even worst for him of all his alternatives, though he ought to do it. Perhaps this second constraint is also obvious.

But even if it is obvious, it does have a certain bite, for there are views according to which it is mistaken: on those views, it can't be, can't really be, can't really be *in the long run* or *all things considered* good for a person to do what he ought not, or bad for him to do what he ought to do. Friends of these ideas think that moral requirement in particular has this feature. On their view, the nature of a person, and the content of moral requirement, are such that necessarily, a person's obedience to the requirement conduces to the good of the person. This strikes me as an excessively high-minded conception of the nature of a person, and therefore of what is good for a person; I will not argue against it, but merely set it aside.

The third constraint is this: the theory must allow for the possibility that a person's doing a thing would not be good for him, even though it is what he most wants to do—and indeed, even if his wants are appropriately restricted.

What I have in mind here is this. There is something right about the idea that what is good for a person is importantly connected with satisfaction of his wants. If a man most wants to become a veterinarian, what could possibly make it fail to be good for him to do so? No doubt he'd make more money as a lawyer, but that hardly matters. If a woman most wants to spend her free time knitting scarves for her friends, what could possibly make it fail to be good for her to do so? On some views, then, what is good for a person is not just importantly connected with satisfaction of his wants, but analyzable into satisfaction of wants—thus: what is good for a person is what conduces to satisfaction of his wants.

But no one thinks this idea even remotely plausible unless the wants that matter to goodness for a person are appropriately restricted. The relevant wants are not passing wants, as where a person has a momentary desire to kick somebody in the shin. Nor wants due to ignorance, as where a person wants to drink a certain liquid, not knowing that it con-

tains cyanide. Nor wants due to improper preference-bending, as where a person has been hypnotized. The relevant wants are, rather, relatively stable, resting on correct information, and autonomously arrived at. Another way to put the point is this: the relevant wants are those a person would have “in a cool hour,” in possession of full relevant information, and under no improper pressure to conform his will to that of others. These are vague conditions, but it is an intuitively attractive idea that assessing what would be good for a person does require assessing exactly these things, hard though it may be to arrive at a correct judgment about them.

Analogously for badness for a person: what is bad for a person is what interferes with satisfaction of his relevant wants.⁴

Now the third constraint I impose on a theory of goodness for people rules this attractive idea out, so it has even more bite than the second constraint. My reason for thinking the idea should be ruled out is this. Consider a doctor who leaves her comfortable practice to go across country to help deal with an epidemic at risk of her own health, perhaps even of her life. Going across country to do this is what she most wants to do, and we can suppose that her wants meet the restrictions I mentioned; are we to say that her doing so is good for her? Or again, consider a man who gives a kidney to his friend, who needs it for life, despite the risk to his health of doing so. Giving the kidney is what he most wants to do, and we can suppose that his wants too meet the restrictions I mentioned; are we to say here that his doing so is good for him? Surely it is intuitively wrong to say that these acts are good for their agents. And we make too little of the moral impressiveness of what they do if we say that what they do is good for them. Their acts are morally impressive precisely because, though good for others, they are not good for them.

⁴This account of goodness and badness for a person is complicated. I think it pays to draw attention to the fact that there is no future in the idea that goodness for a person is analyzable, more simply, in terms of needs. A person needs food and drink, and it is good for the person to get these things. But your doing a thing might be good for you without its being the case that you *need* it done. For example, it might well be good for you to learn Russian without its being the case that you need to learn Russian. What a person needs is not just anything that it would be good for him to get, but rather what it would be bad for him to not get. And I see no more reason to think (indeed, less reason to think) that badness for a person is analyzable in terms of needs than there is to think that needs are analyzable in terms of badness-for.

So also for animals, plants, and artifacts. Animals and plants also need food and drink, and it is good for them to get these things. A lawn mower may need oiling; if so, it would be good for it to get oiled. But the fact that animals, plants, and artifacts need these things is a product of its being bad for them to not get them.

The possibility I point to here is not restricted to cases in which an act is morally impressive. Consider a master chess player who spends his time studying chess, gets no exercise, and smokes heavily because he finds that smoking helps him to concentrate. These are things he wants to do because he wants to become world champion. But I think we cannot at all plausibly say that doing them is good for him.⁵

What works against the idea that what is good for a person is what conduces to satisfaction of his wants is the idea that what is good for a person is what conduces to his health. Many people do sometimes want most to do what interferes with their health, and this even though their wants do meet the restrictions I mentioned above.

So perhaps we should instead opt for a theory according to which goodness for people is analyzable into conduciveness to health? There are at least two objections to this idea. First, the idea is far too narrow. It might be good for your daughter to be accepted by a good law school; I find it hard to believe that this would conduce to her health. Second, and more interesting, if we ask why it is good for a person to be healthy, there seems to be an answer: health does of course have its special pleasures, but what seems fundamental to its value to us is that it is a prerequisite for our being able to do much of what most people most deeply want to do. These two objections head us back toward the idea that goodness for people consists in satisfaction of wants.

My own impression is that the best theory of what is good for a person lies somewhere in between these two ideas. If this is right, then it is clear that, if not exactly how, a person's wants are relevant to what he ought to do: they are so by being relevant to what it would be good for him to do.

Fortunately for our purposes, it does not really matter exactly where, between those two ideas, the best theory of what is good for a person lies. Our main concern will be what is to be said of cases in which a person's doing a thing would be good (or bad) for him, but bad (or good) for another person. So long as the criteria for goodness for and badness for a person are general, in the sense that they bear on both parties—and so

⁵ This is obvious enough if his efforts do not succeed, so that he does not become champion. But what if they do and he therefore does? Should we say (i) that his becoming champion is good for him, though his taking the necessary means to that end was not? Or (ii) that, given those were the necessary means to that end, and that they were not good for him, it follows that his becoming champion is not itself good for him? or (iii) that, given his becoming champion is good for him, it follows that his taking those means was, after all, good for him? I leave this open.

long as the three constraints I listed are met—that will suit us well enough.

4.

Less has been written about what is good or bad for an animal or plant, and hardly anything about what is good or bad for an artifact or a river. I will have to be brief about these matters.⁶

What is good for a plant is obviously what conduces to its health. I said there were two objections to the idea that goodness for people is analyzable into conduciveness to health; neither arises in the case of plants. First, everything that is good for a plant, whatever it may be, conduces to its health: more strongly, if something is good for a plant, then that is because and only because it conduces to the plant's health. Second, if we ask why it is good for a plant to be healthy, there is no answer. Plants have no wants, and a fortiori, their being healthy is not good for them because it is a prerequisite for their being able to do what they want.

This difference between what is good for people and what is good for plants is pretty plainly due to the very fact that people have wants and plants do not.

Even the higher animals are more like plants in this respect than they are like people. Unlike plants, they do have wants; unlike people, their range of wants is narrow. (No doubt that is due to the fact that their range of experiences is narrow. They have a more or less wide variety of sensory experiences, and they feel pleasure and pain; but they feel no resentment, indignation, ambition, pride, envy, admiration, and so on.) In light of the fact that they are conscious beings, a plausible account of what is good for them should presumably lie somewhere between that for people and that for plants.

What is it for a person, animal, or plant to be healthy? Curiously enough, I think we are helped if we have a look at artifacts first.

Your doing a thing might be good for your lawn mower—for example, oiling it might be good for it. How so? Oiling it would conduce to its being in good condition. What does that come to? Oiling it would

⁶My own views about these matters appear, in greater detail, in "The Right and the Good," *Journal of Philosophy* 94, no. 6. Though I still think that my remarks there about goodness for other things are true, my remarks there about goodness for people now strike me as false—they are argued against in section 3 above.

conduce to its being able to do what it was designed by its manufacturer to do, namely enable a person who uses it properly to mow lawns easily, safely, and well. We don't call an artifact that is in good condition a healthy artifact, but the idea is surely similar.

Similar in that a living creature is healthy just in case it is able to do what it was 'designed by nature' to do—among other things, grow in ways suitable to members of its species, and propagate itself. Animals are designed by nature to do more than plants are, and human beings to do more things still. But feeding any living creature the food suitable to members of its species, and in a suitable amount, would be good for it in that doing so would conduce to its being in a condition in which it is able to do those things.

It is important to recognize that what is good for an animal or plant is not reducible to what is good for people. It might be good for a cat's owner to have it declawed; it is nevertheless not good for the cat to do so. It might be good for a lawn's owner to kill the dandelions that grow in it; it is not good for the dandelions to do so. What is good for human beings plays a more subtle role in the case of artifacts. There are such things as artifacts because of human wants. But whether or not doing a thing is good for an artifact is not a function of human wants. I might want my lawn mower to become unusable to mow lawns with; that is compatible with its being bad for the lawn mower to do so. So also, I should think, for animals and plants that are not in fact designed by nature but instead by human beings—that is, animals and plants that are bred by humans for special purposes: once bred, what is good for them is a function of their design and not our purposes.

Rivers, lakes, mountains, and ravines are quite another matter. In their case, it is what is good for human beings that fixes what is good for them. The Charles River in Massachusetts had become polluted; it was good for the river that measures be taken to decrease the pollution. But that was not because doing so conduced to the river's being able to do what rivers are designed by nature to do: there is no such thing. Rather it was because human beings wished to be able to sail, fish, and swim in the river, and to avoid the nasty look and smell that pollution brings with it.

These differences must surely have a bearing on whether a person ought or ought not or may do a thing where the person's doing it would be best for him but bad for a thing of one or another of these kinds. A theory about what a person ought to do should deal with them. I have

no such theory, however, and will only be able to indicate the structure into which I suggest they should be fitted.

5.

What we have been looking at so far are the easy cases, namely those in which

(Isolation) *A*'s *V*-ing would neither be bad for anything other than *A* nor infringe anything's rights

is true. The first of my four suggestions was about those easy cases, and it said:

(I) If (Isolation) is true, and if also *A*'s *V*-ing would be better for him than his doing any of the other things it is open to him to do instead, then it follows that *A* ought to *V*.

We should turn now to the hard cases, namely those in which (Isolation) is false. They are the ones of most concern to moral philosophy.

The second of my four suggestions is this:

(II) Justice and generosity are second-order ways of being good; injustice and (what I will call) miserliness are second-order ways of being bad.

For a reason that will emerge in section 7, it is injustice and miserliness that will primarily matter to us. So let us begin with injustice.

To act unjustly is to fail to do what is owed; that is, it is to fail to do what another person has a right to one's doing. More precisely, to act unjustly is to fail to do what another person has a *non-overridden* right to one's doing. Suppose that Alfred has a right that you meet him at the corner of High and Main at 3:00 P.M. today. How might Alfred have acquired that right? Perhaps you promised him that you would meet him there then. Here are two kinds of case in which Alfred's right is overridden. (i) Bertha has a more stringent right that you not meet Alfred there then. Perhaps you promised her that you would not do so, which is what gave her a right that you not do so. What would mark her right as more stringent? Well, perhaps while Alfred will merely be annoyed if you do not turn up on time, she will suffer greatly if you do. (ii) Alfred's right that you meet him there then is weak (that is, nonstringent), and no one

has a right that you not be there then, but your being there then would be very bad for you, or for Bertha, or for Bertha and Charles. Perhaps while Alfred will (once again) merely be annoyed if you do not turn up on time, you or Bertha or both Bertha and Charles will suffer greatly if you do. In cases of kinds (i) and (ii), Alfred's right that you meet him at the corner of High and Main at 3:00 P.M. today is overridden; and if you fail to meet him there then, your doing so is not unjust.⁷

By hypothesis, Alfred has a right that you meet him at the corner of High and Main at 3:00 P.M. today. What if (iii) Bertha has an equally stringent right that you not meet him there then? (For simplicity, let us suppose that no one else would be affected, whichever you do.) Here we have a tie again. I think that nothing theoretically interesting turns on which choice we make; since I think it simpler, and somewhat more plausible, to conclude that neither of them has a non-overridden right in respect of your meeting or not meeting Alfred, I will assume so. (A good idea for you in such a situation would be to try to get them to agree to your flipping a coin. If they do, then presumably the winner, and only the winner, has a non-overridden right.)

In sum, then: to act unjustly is to fail to do what something has a non-overridden right to one's doing.

A feature of this account of acting unjustly calls for explicit mention. What I refer to is the fact that according to this account, the question whether a person's act is unjust does not turn on what the person's beliefs or intentions are. What fixes whether the person's act is unjust is not the person's subjective state of mind, but instead the objective fact that something else does or does not have a non-overridden right against him. That seems to me exactly as it should be. Suppose Alfred does not believe that anything has a non-overridden right that he send Bertha a check for fifty dollars. But suppose that is because he has sim-

⁷Some philosophers hold the view that there is no such thing as a right that is overridden; on their view, all rights are 'absolute.' Thus they would say that in cases of kinds (i) and (ii), Alfred has no right that you be there then, for the facts about Bertha in (i) and about you or Bertha or Bertha and Charles in (ii) make Alfred cease to have the right you gave him when you made your promise to him. Other philosophers—I include myself—hold the view that Alfred does have the right you gave him, though it is overridden, and justice therefore does not require your according it to him. My reason for preferring this second view is the fact that if you do not keep your promise to Alfred, you will at a minimum owe him an apology: this fact seems to me to show that the right you gave him did not simply go out of existence. I will throughout assume without argument that this second view is correct. For our purposes, it does not really matter which is correct, and what I say below could easily enough be revised in such a way as to square with the first view.

I supply more detail on these matters in *The Realm of Rights* (Cambridge: Harvard University Press, 1990).

ply forgotten that he had promised her that he would, and suppose also that her right that he do so is non-overridden. Then his failing to send her a check would be unjust. If it is not his fault that he forgot his promise, then we may well think he would not be to blame for failing to send her a check; but that his failure would be unjust seems quite clear. I will summarize this point by saying that the question whether a person's act is unjust is objective—just as, as I said earlier, the question whether a person ought to do a thing is objective. I think it clear that this is true of justice too.⁸

Let us turn now to miserliness. I will so use that word that it stands for the contrary of generosity. My use of the word is therefore broader than the ordinary use of it, for I think that we call an act miserly only where its agent is mean or grudging about money. Suppose that Bertha is drowning and Alfred alone can save her, as it might be, by throwing her his life preserver. If Alfred refuses to do this on the ground that he is feeling tired, and does not want to bother, then I will say that his refusal to throw it is miserly, despite the fact that his refusal is not due to meanness about money.

By way of preliminary characterization, I will describe as “miserly” any act that provides its agent with a small gain—whether financial or otherwise—or even no gain at all, despite the act's being very bad for another person. Thus an act of gratuitous cruelty, inconsiderateness, or discourtesy may also be miserly.

Why “another person”? Is it only where the ‘victim’ of the act is a person that an act is miserly? I will suppose not. If animals lack rights, then it is not possible for one's treatment of them to be unjust; I mean to allow that it is possible for one's treatment of an animal to be miserly whether or not animals have rights. This use of “miserly” is as it should be if miserliness is to be the contrary of generosity, for I should think it possible to be generous to an animal.

What about plants, artifacts, and rivers? Rivers we can set aside, since nothing is bad for them except by way of being bad for people. I am inclined to think, however, that we should allow for the possibility that one's treatment of a plant or artifact is miserly. If you gain nothing

⁸It might pay to make explicit that my point in the text has to do with *acts* only. The questions whether a *person* is unjust and whether a person is just are not, or anyway are not wholly, objective. I will shortly opt for a similar point about the questions whether an act is miserly, and whether an act is generous; opting for it is compatible with accepting that the questions whether a person is miserly, and whether a person is generous, are not, or anyway are not wholly, objective.

from destroying a tree or a painting, by which I mean to include that it would cost you nothing to refrain from destroying it, then I will say that your treatment of it is miserly—even if there is no person whose property rights you violate, or whom you affect for the worse.

So here is an interesting difference between injustice and miserliness: your treatment of a thing can be unjust only if the thing is a right-holder, whereas your treatment of a thing can be miserly even if it is not. I will draw attention to yet another difference between injustice and miserliness shortly.

In any case, let us emend the preliminary characterization of miserliness that I gave above. It is surely right to think that Alfred's *V*-ing is miserly only if

- (i) Alfred's *V*-ing provides Alfred with at most a small gain, despite its being very bad for something else

is true.

Another condition is necessary, however. Suppose that (i) is true of Alfred's *V*-ing. Suppose, however, that someone has a non-overridden right to Alfred's *V*-ing. Then I will say that Alfred's *V*-ing is not miserly.

Here is an example. Suppose that Bertha is using Alfred's typewriter without his permission; she needs it to complete some work she is doing, the doing of which is important to her. Suppose that Alfred has no need of the typewriter himself, and indeed, that he would gain nothing at all by requesting its return. Suppose he nevertheless requests its return. Then his requesting its return gains him nothing, despite being very bad for Bertha. It sounds, therefore, as if Alfred's requesting the return of the typewriter is miserly. But suppose, last, that Cora has a non-overridden right that Alfred request the return of the typewriter. (Alfred has promised Cora that he will get it back from Bertha to give to Cora, who needs it more than Bertha does.) Now it is clear that Alfred's requesting the return of the typewriter is not miserly.

So we need to add a further condition on miserliness: Alfred's *V*-ing is miserly only if

- (ii) Nobody has a non-overridden right that Alfred *V*

is also true.

Condition (ii) is not unique to miserliness, for it holds of generosity too. If somebody has a non-overridden right that Alfred *V*, then his *V*-

ing is not miserly. So similarly: if somebody has a non-overridden right that Alfred V, then his V-ing is not generous. (It is not generous in a person to pay his grocer's bill on time, however good for his grocer his doing so would be.)

In sum, I take it that conditions (i) and (ii) are both necessary for an act to be miserly. Are they jointly sufficient? I should think not. That is because of a feature of miserliness—and generosity—that I will call context-dependence.

Suppose that I have never given to the needy, and it is proposed to me now that I make a small contribution to them. If I refuse, my refusing is miserly; indeed, it is very miserly. By contrast, suppose that you have very often given to the needy, and it is proposed to you now that you make yet another small contribution to them. If you refuse, your refusal is markedly less miserly, and perhaps is not miserly at all. (Why so? Presumably because your total contribution is already not small.) This feature of miserliness is what I am calling context-dependence.

Context-dependence is also a feature of generosity: your making a further small contribution now is the more generous given you have very often contributed in the past.

But context-dependence is not a feature of either injustice or justice, and therefore constitutes yet another difference between injustice and miserliness. My failing to pay my grocer's bill on time is not less unjust if I have very often paid my grocer's bill on time in the past. Nor is my paying my grocer's bill on time more just if I have paid it on time very often in the past.⁹

Expressing the context-dependence of miserliness in the form of a further condition on miserliness would be a complex business, and not, I think, an interesting one. So I leave open how this is to be done. I leave open also the possibility that further conditions should be placed on miserliness.

But what I do not leave open is the possibility of adding a condition concerning the agent's subjective state of mind. I will so use the word "miserly" that the answer to the question whether a person's act is or is not miserly does not turn on his beliefs or intentions: rather it turns on

⁹I will suggest in section 7 below that we ought to avoid miserliness and injustice. Taking liberties with Kant's terminology, we might say that we have perfect duties to avoid miserliness and injustice in the sense that we are on all occasions to avoid both—but that we also have an imperfect duty to avoid miserliness in the sense that miserliness (unlike injustice) is context-dependent.

what he (objectively) does, as condition (i) requires, and on whether nobody has a non-overridden right against him, as condition (ii) requires. I will summarize this point by saying that the question whether an act is miserly is objective. Miserliness, however, was to be the contrary of generosity; and isn't our ordinary use of "generous" subjective? I think it is, or that it at least mostly is. Suppose that muddled Alfred intends to be helping Bertha, and believes that he is succeeding in doing so, but he is mistaken: what he is doing in fact harms her. Do we say that his act was nevertheless generous? I think we do.

On the other hand, I am not sure that our use of "generous" is always subjective in this way. Consider the muddled miser, who intends not to help, and thinks he is not helping. Unbeknownst to him, however, he is in fact helping. Mightn't we say that unbeknownst to him, his act was generous? (If a person says "I didn't mean to be generous," do we always conclude that his act was not generous? Mightn't we sometimes reply "All the same, you were"?)

In any case, I will throughout use the word "generous" in such a way that it—like its intended contrary, "miserly"—is objective. The muddled miser, whose act was (as I will say) generous certainly deserves no praise for it; he may even deserve blame for it. By contrast, while muddled Alfred's act was (as I will say) not generous, and perhaps was even miserly, he may deserve no blame for it; he may even deserve praise for it.

6.

To return now to my suggestions. The second of my four suggestions was:

(II) Justice and generosity are second-order ways of being good; injustice and (what I will call) miserliness are second-order ways of being bad.

We now have in hand an account of injustice and anyway a good enough preliminary account of miserliness to be able to see, as follows, that (II) is true.

It is presumably clear enough that injustice and miserliness are ways of being bad, and justice and generosity are ways of being good.

My ground for saying that injustice is a second-order way of being

bad has already emerged, for what I have in mind about it is this. The answer to the question whether an act is unjust turns in part on whether those affected by it have the relevant rights; but it also turns in part on whether those rights are overridden, and that turns in part on how good or bad the act would be for those affected by it. Goodness-for and badness-for are first-order ways of being good; whether an act is unjust cannot be settled unless it is first settled how good or bad it would be in the relevant first-order ways.

That miserliness is a second-order way of being bad is perhaps even clearer than that injustice is—for it is clear that whether an act is miserly is in part a function of how good refraining would have been for others, and how bad refraining would have been for the agent.

The analogue of this point also holds of justice and generosity.

A word or two should probably be said here about moral goodness and badness. I said that injustice and miserliness are ways of being bad, and justice and generosity are ways of being good. Couldn't we have said that the first pair are ways of being morally bad, and the second pair ways of being morally good?

I am inclined to think that the ordinary use of the terms "morally good" and "morally bad" is such that the questions whether they apply to an act are subjective, not objective. Thus consider again a person who means to be acting justly, and thinks he is succeeding in doing so. As I said, he may be mistaken: it may be that his act is in fact unjust. If the mistake is not his fault, then (as I said) we may well think he is not to blame for his unjust act. Indeed, he might even deserve praise for it. Suppose he does deserve praise for it. Then I think we would conclude that, although his act was unjust, it was morally good.

Similarly, consider again the muddled miser, who intends not to help, and thinks he is not helping. Unbeknownst to him, however, he is in fact helping. I said I would so use the word "generous" that his act is generous. He does not of course deserve praise for it; rather, he may deserve blame for it. Suppose he does deserve blame for it. Then I think we would conclude that, although his act was generous, it was morally bad.

In short, our ordinary use of those terms seems to me to be such that their applicability to an act turns on its agent's subjective state of mind, in particular, on his beliefs and intentions.

I see no good reason to reject that usage, and I therefore think we should agree that the questions whether an act is morally good or

morally bad cut across the questions whether it is just or unjust, generous or miserly.

Moreover, we should agree that the questions whether an act is morally good or morally bad cuts across the questions whether its agent ought or ought not be doing what he is doing—for, as I argued earlier, the questions whether a person ought to do a thing, or ought not do it, are not subjective, but instead objective.

Our concern here is what a person ought or ought not do. I therefore bypass moral goodness and badness.

7.

The third of my four suggestions is about what a person ought to do:

(III) *A* ought to *V* if his not *V*-ing would be either unjust or miserly, and *A* ought not *V* if his *V*-ing would be either unjust or miserly.

I think that, given the characterizations of injustice and miserliness I offered in section 5, this suggestion is intuitively very plausible.

We should notice that (III) says “if,” and not “if and only if.” That is because there may well be other grounds for saying that a person ought to do a thing or ought not do it than injustice and miserliness. For example, my first suggestion said:

(I) If (Isolation) is true, and if also *A*'s *V*-ing would be better for him than his doing any of the other things it is open to him to do instead, then it follows that *A* ought to *V*—

where (Isolation) says:

(Isolation) *A*'s *V*-ing would neither be bad for anything other than *A*, nor infringe anything's rights.

Wherever (Isolation) is true, *A*'s *V*-ing is neither unjust nor miserly; yet it may all the same be the case that *A* ought to *V*, or ought to refrain from *V*-ing. And we should not suppose that the word “ought” has different senses, according as injustice and miserliness are or are not in question. I have invited you to agree that there is such a thing as *the* advice sense of the word “ought,” which is the sense that I am taking it to have throughout.

It is worth noticing that, given (III),

(II) Justice and generosity are second-order ways of being good; injustice and (what I will call) miserliness are second-order ways of being bad

is more important than it may initially have appeared. Suppose that Alfred's doing a thing would be good for him, but that his doing it would be unjust. If (III) is true, then Alfred ought not do the thing. That, I am sure, will seem plausible to most people. Perhaps not to others. But whether it seems plausible or not, many people ask for a justification. *Why* does the fact that Alfred's doing the thing would be unjust yield that he ought not do it?—when, by hypothesis, it would be good for him to do it?

Another way to put this question is this. I said in Part I that the fact that a person's doing a thing would be good in a way counts in favor of his doing it, and the fact that a person's doing a thing would be bad in a way counts against his doing it. For (III) to be true, it is required that the fact that Alfred's doing a thing would be unjust (hence bad in a way) counts more strongly against his doing it than the fact that his doing it would be good for him counts in favor of his doing it. Why (it is asked) should we think this true?

If (II) is true, this question has not the interest it has been thought to have. According to (II), we are not to weigh the fact that Alfred's doing the thing would be good for him against the fact that his doing it would be unjust, leaving open the possibility that the fact that his doing it would be good for him counts more strongly in favor of his doing it. For given that injustice is a second-order way of being bad, the fact that his doing the thing would be good for him has *already* been taken into consideration in arriving at the conclusion that his doing it would be unjust. To allow the fact that his doing the thing would be good for him to have a further, independent, bearing on what he ought to do would be to double-count it.

There is a more general point in the offing. Suppose that Alfred's doing a thing would be unjust or miserly. According to (III), it follows that Alfred ought not do it. In saying that Alfred ought not do it, I am saying that Alfred ought not do it, whatever the consequences of his doing it—that is, however good the consequences of his doing it may be in

whatever ways they may be good. That his act would possess the second-order features of being unjust or miserly is a conclusion arrivable at only after consideration of the first-order ways in which Alfred's doing the thing would be good.

Suggestions (I), (II), and (III) can be connected in the following way: we can say that they jointly yield that what a person ought to do is what gives adequate weight to the interests of all who would be affected. Suppose that (Isolation) is true of Alfred's drinking some hot lemonade. It follows that only Alfred's own interests are relevant to the question whether he ought to drink some. According to (I), he ought to drink some if his doing so is best for him. And we can say: his doing so would give adequate weight to his own interests. His choosing any option less good for him would give inadequate weight to his own interests—it would be imprudent.

Suppose, however, that (Isolation) is not true of Alfred's drinking some hot lemonade. (Perhaps he has promised Bertha that he would not.) Then the interests of others are relevant to the question whether he ought to drink some. According to (III), he ought not drink some if his doing so would be unjust or miserly; according to (II), the question whether it would be turns in part on how good or bad his doing so would be for him and for those others. If his doing so would be unjust or miserly, so that he ought not, then we can say: his doing so would be his failing to give adequate weight to the interests of all who would be affected—those of the others as well as his own.

8.

But is

(III) *A* ought to *V* if his not *V*-ing would be either unjust or miserly, and *A* ought not *V* if his *V*-ing would be either unjust or miserly

true? There is a possible objection to it that I think worth taking note of.

Suppose that *A*'s *V*-ing would be unjust. According to (III), it follows that *A* ought not *V*. Can we consistently also suppose that his not *V*-ing would be miserly? Suppose we can. Let us do so, then: we now suppose also that *A*'s not *V*-ing would be miserly. According to (III), it follows that he ought to *V*. So it follows both that *A* ought not *V*, and also that *A* ought to *V*.

Some philosophers are quite content to make room for the possibility that

(1) *A ought not V*

and

(2) *A ought to V*

are both true.¹⁰

Other philosophers—I include myself—think it is not possible that (1) and (2) are both true. We must therefore either reject (III) or rebut the objection to it that I have just drawn attention to.

The rebuttal is easy, however, given my characterizations of injustice and miserliness. I said just above: suppose that *A*'s *V*-ing would be unjust. I then asked: can we consistently also suppose that his not *V*-ing would be miserly? The answer is that we can't. Given my characterization of injustice, *A*'s *V*-ing is unjust only if

(3) Somebody has a non-overridden right that *A* not *V*

is true. But given my characterization of miserliness, *A*'s not *V*-ing is miserly only if

(4) Nobody has a non-overridden right that *A* not *V*

is true. It is plain that (3) and (4) are incompatible. So it cannot be the case both that Alfred's *V*-ing is unjust and that his not *V*-ing is miserly.

9.

My fourth suggestion emerges on consideration of the question why I have focused on injustice and miserliness, which are ways of being bad, rather than on justice and generosity, which are ways of being good.

To begin with justice. According to (III), the fact that *A*'s *V*-ing

¹⁰It is obviously not possible that (2) and

(*) It is not the case that *A* ought to *V* are both true, so that if (1) entails (*), then those philosophers are mistaken—that is, it is not possible that (1) and (2) are both true. But those philosophers simply deny that (1) entails (*).

The philosophers I refer to here are those who accept that there are what they call "moral dilemmas." I indicated earlier that my own use of "ought" would be all-things-considered, and it is surely clear that if "ought" is construed as I construe it, then (1) does entail (*), and (1) is therefore incompatible with (2).

would be unjust entails that *A* ought not *V*. But the fact that *A*'s *V*-ing would be just does not entail that *A* ought to *V*. That, I should think, is obvious. There might well be cases in which you have several options for action—as it might be, several ways of distributing a benefit you owe—each entirely just, there being no one of them that you ought to choose.

A similar point holds of generosity. It might well be that while you cannot give aid to both Alfred and Bertha, giving aid to Alfred and giving aid to Bertha would each be generous, neither option being such that you ought to choose it.

In short, we cannot say that a person ought to do a thing if his doing of it would be either just or generous.

A further, and more interesting, reason why we cannot say that a person ought to do a thing if his doing of it would be either just or generous issues from a fact about generosity. An act can be more or less generous, and a very generous act may be, as it is sometimes put, supererogatory, that is, as we say, 'above and beyond the call of duty.' The person who saved a life at risk of his own acted very generously. It does not follow that he did what he ought. Indeed, it would not have been true to say of him that he ought to save that life. What he did for the other person was beyond what he ought.

This point does not also hold of justice. While some generous acts are beyond what one ought, no just act is. No matter, for our purposes. The fact remains that we cannot say that a person ought to do a thing if his doing of it would be either just *or* generous.

I said earlier that what a person ought to do is what gives adequate weight to the interests of all affected—not lavish weight. What "You ought" requires of us is the adequate minimum. That is my fourth suggestion:

(IV) Doing what one ought only requires giving the adequate minimum weight to the interests of all who are affected.

It pays to stress, however, that accepting (IV) does not commit us to a conclusion about praise. Certainly the person who saved a life at risk of his own deserves much praise for doing so: and I should think, by contrast, that one rarely deserves praise for doing what one ought. But it may on occasion be the case that a person does deserve praise for doing what he ought—as, for example, where his avoiding injustice requires more courage and a stronger will than are normally required in our dealings with others.

Suggestion (IV) says *all* who are affected, which of course includes *A* himself. On the one hand, whether someone's right against *A* is non-overridden turns on what is good or bad for *A* as well as on what is good or bad for others. On the other hand, there are limits to what avoiding miserliness requires of us. I said that the person who saved a life at risk of his own acted very generously, and deserves much praise for doing so. Would it be very generous, deserving of much praise, to cure someone's sore throat at risk of one's life? I hardly think so. The beneficiary's gain is too small in comparison with the loss the agent risks. There is nothing praiseworthy in making a major sacrifice in order to forestall another person's small loss. Surely such an act is not merely not praiseworthy, its agent ought not perform it. Why so? Presumably the answer lies in the fact that a person who makes a major sacrifice in order to forestall another person's small loss does not give adequate weight to his own interests—and a fortiori does not give adequate weight to the interests of all who are affected.

10.

It is time now to make contact with the proposals I made in Part I about reasons for action. I said there that the concept 'reason for a person to do such and such' has been thought to have an intimate connection with the concept 'ought'. Some philosophers hold that *A* ought to *V* only if there is a reason for *A* to *V*. Or even more strongly: *A* ought to *V* if and only if *V*-ing is precisely what there is most reason for *A* to do.

I then proposed that a fact is a reason for *A* to *V* if and only if it is a fact to the effect that

A's *V*-ing would be good in a way, or
A's not *V*-ing would be bad in a way, or
A's *V*-ing would be better in a way than *A*'s doing anything else, or
Someone has a right that *A V*, or . . . ,

leaving open that other evaluative facts should be added to this list.

What has emerged in the preceding sections is the idea that some of these kinds of fact are good reason to believe that *A* ought to *V*. Indeed, some are themselves conclusive reasons to believe that *A* ought to *V*. Thus the fact that *A*'s not *V*-ing would be unjust, or would be miserly, is conclusive reason to believe that *A* ought to *V*. And some are—not

themselves, but in conjunction with other facts—conclusive reasons to believe that *A* ought to *V*. Thus the fact that *A*'s *V*-ing would be better for him than his doing anything else is not itself conclusive reason to believe *A* ought to *V*, but the conjunction of it with (Isolation) is.

Suppose we say that wherever there is a fact that is—or is, in conjunction with (Isolation)—conclusive reason to believe that *A* ought to *V*, then there is “most reason” for *A* to *V*. Then we can say: *A* ought to *V* if *V*-ing is what there is most reason for *A* to do.

Can we also say: *A* ought to *V* only if *V*-ing is what there is most reason for *A* to do? I should think it very plausible that we can. That a person ought to do a thing is not a fact that floats free of anything that might be thought to make it a fact; and what would make it a fact is surely facts of the kinds we have been looking at. Have we have looked at all the relevant kinds? An agent's advantage, where (Isolation) is true, covers a good bit of territory. So does injustice. So does miserliness, understood in the way I have been taking it. (As I said, an act of gratuitous cruelty, inconsiderateness, or discourtesy may be miserly.) But I will not try to argue that there are no others.

In any case, if the stronger claim—namely that *A* ought to *V* if and only if *V*-ing is what there is most reason for *A* to do—is true, then so also is the weaker claim that *A* ought to *V* only if there is reason for *A* to *V*.

However, the availability of these conclusions turns on our opting for the broad construal of reasons for action. According to the narrow construal, a fact to the effect that *A*'s *V*-ing would be good in a way is a reason for *A* to *V* only if *A* wants to do something good in that way. Similarly, a fact to the effect that *A*'s not-*V*-ing would be bad in a way is a reason for *A* to *V* only if *A* wants to avoid doing something bad in that way. For example, the fact that *A*'s not-*V*-ing would be unjust or miserly is a reason for *A* to *V* only if *A* wants to avoid injustice or miserliness. As I said in Part I, section 13, this narrow construal of reasons for action strikes me as out of accord with our ideas about reasons for action.

But as I also said, I see no theoretically important reason for rejecting this narrow construal. Or for opting for it. Whatever we wish to accomplish in moral theory had better be accomplishable whether we opt for this narrow construal of reasons for action or for the broad one. There is something interesting that motivates opting for the narrow one, and I will discuss it in the following section. It is enough for our purposes to notice here that whether we call the kinds of facts I pointed to above “reasons for *A* to *V*” does not really matter. Some of them are—or to-

gether with (Isolation) are—conclusive reasons for believing that *A* ought to *V*. *That* is the point important to moral theory.

II.

Why is it that many contemporary moral philosophers have focused on the concept ‘reason for action’? Why do many of them think, in particular, that *A* ought to *V* if and only if *V*-ing is what there is most reason for *A* to do? I suggest that that is because they wish to have it turn out that it is always *rational* for a person to do what he ought. Indeed, that rationality requires doing what one ought.

They swim upstream, however. That is because there is a familiar, because attractive, theory of rationality in action according to which *A*’s *V*-ing would be rational if and only if it would satisfy *A*’s wants. More precisely: if and only if it would satisfy *A*’s appropriately restricted wants. As I said in section 3, the restriction excludes merely passing wants, wants due to ignorance, and wants due to improper preference-bending. As I also said in section 3, the fact that *A*’s *V*-ing would satisfy *A*’s appropriately restricted wants cannot be thought sufficient to mark *A*’s *V*-ing as good for him. But it is no surprise that many people suppose that fact is sufficient to mark *A*’s *V*-ing as rational.

If one accepts this familiar theory of rationality in action, then one cannot plausibly also accept that it is always rational for a person to do what he ought. For it is wildly implausible to suppose that wherever a person ought to do a thing, it will also be the case that his doing it would satisfy his wants—even if we appropriately restrict his wants.

So those many contemporary moral philosophers who think that *A* ought to *V* if and only if *V*-ing is what there is most reason for *A* to do—and think this because they wish to have it turn out that it is always rational for a person to do what he ought—have a hard job ahead of them. They need to find an alternative to that familiar theory of rationality in action.

Its very attractiveness is in fact what motivates opting for the narrow construal of reasons for action, according to which a fact is not a reason for action unless the agent has the appropriate wants—thus, for example, that the fact that *A*’s *V*-ing would be good in a certain way is a reason for *A* to *V* only if *A* wants to do something good in that way. If one accepts that familiar theory of rationality in action, then one is very

likely to think that a fact is a reason for *A* to *V* only if *A*'s acting on it would conduce to the satisfying of *A*'s wants. Moral philosophers who think that *A* ought to *V* if and only if *V*-ing is what there is most reason for *A* to do need to be able to reject the narrow construal of reasons for action; so (to repeat) they need to find an alternative to that familiar theory of rationality in action.

There surely is some pressure on us, whatever our views about reasons for action, to find an alternative. After all, it would be at a minimum unfortunate to have to agree that Alfred's paying his grocer's bill is irrational if his wants would be better satisfied if he did not. Moreover, we give advice when we say "You ought"; how can it be thought coherent to advise Alfred to pay his grocer's bill in the words "You ought to pay it, though I grant that your paying it would be irrational"?

But what alternative is available? One possibility begins by fixing on the fact that we are taking "ought" to have only one advice sense. We are supposing that its one advice sense is 'all things considered,' where among the things considered is what would be good or bad for the agent. And it might be asked: "How could it be rational to believe that you ought (in that sense) to do a thing, and yet not do it?"¹¹ One obvious objection to the idea that this couldn't be rational is rebuttable. Thus suppose someone objects that a person might perfectly well believe he ought to do a thing and yet not want to do it, and therefore not do it. A friend of this idea replies: "I don't deny that this is possible, I say only that it is irrational."

Curiously enough, support for the idea can be found in a phenomenon observable in many of those who accept the familiar theory of rationality in action. What I refer to is the fact that they are inclined to resist the idea that a person ought to do a thing when, as it turns out, it would be (on the familiar theory) irrational for him to do it. They are inclined to say that if it would be irrational for the person to do the thing, then it is at a minimum doubtful that he ought to do it. *Nobody* is happy to allow that it might be true to say to Alfred: "You ought to pay your grocer's bill, though your paying it would be irrational."

My own impression, however, is that it just is not clear enough what the dispute between those who accept the familiar theory of rationality in action and those who are in search of an alternative is a dispute *about*. One way in which this unclarity emerges is the following.

Let us look again at reasons for belief. A fact *F* is a reason for believ-

¹¹ I am grateful to Robert Streiffer for drawing my attention to this idea.

ing a hypothesis H just in case F counts in favor of H . For F to count in favor of H is for F to lend weight to H . And there is such a thing as a fact's being a conclusive reason for believing a hypothesis: F is a conclusive reason for believing H just in case F entails H . Suppose that F is a conclusive reason for believing H , and thus that F entails H . Then there is a quite clear sense in which it is 'against reason' for a person to believe that there is such a fact as F , while believing that H is false: it is self-contradictory to believe these things.

A fact F is a reason for A to V just in case F counts in favor of A 's V -ing. I suggested in Part I that for F to count in favor of A 's V -ing is for F to be a fact to the effect that A 's V -ing would be good in a way, or that A 's not V -ing would be bad in a way, or . . . Is there such a thing as a fact's being a conclusive reason for A to V ? Suppose (i) we say that the fact that A ought to V is a conclusive reason for A to V . What could we mean by that? In what sense could it be thought to be 'against reason' for A to fail to V while believing that he ought to? Certainly no fact entails an action; what other sense is available?

I hasten to add that life is no easier for friends of the familiar theory of rationality in action. Suppose (ii) we say that the fact that A 's appropriately restricted wants would be satisfied if he V ed is a conclusive reason for A to V . What could we mean by that? In what sense could it be thought to be 'against reason' for A to fail to V while believing that his appropriately restricted wants would be satisfied if he V -ed?¹²

I think that it would pay to notice something further about the familiar theory of rationality in action. Suppose we accept the equally familiar Humean account of explanation of action. (I drew attention to it in Part I, section 11.) Thus suppose we accept that what a person will in fact do is what he believes will most efficiently satisfy his wants. Suppose, last, that Alfred's wants are all appropriately restricted, and that he believes that refraining from paying his grocer's bill will most efficiently satisfy them. The friend of the familiar theory of rationality says that for Alfred to pay his grocer's bill would be irrational. But that can't be right if Hume is right. For if Hume is right, then Alfred's paying his grocer's bill would not be irrational. It would be unintelligible. It would be inexplicable.

Moreover, if we say to Alfred "You ought to pay your grocer's bill,"

¹²The questions I am raising in the text here are also raised by the (in my view at best suspect) idea that there is such a thing as practical reasoning which contrasts with theoretical reasoning in the following way: while the conclusion of a bit of theoretical reasoning is a proposition, the conclusion of a bit of practical reasoning is an act.

we are not advising him to do something it would be irrational for him to do. If Hume is right, we are instead advising him to do something his doing of which would be inexplicable. So be it. No doubt we do not give a person advice unless we think, or at least hope, that our doing so will affect his wants and thereby his actions. But alas it is not in the least uncommon for efforts of this kind to fail.

Yet isn't the familiar Humean account of explanation of action at least as attractive as the familiar theory of rationality in action?

There is much more to be said on this topic, but I suggest that we bypass it. For my own part, it seems to me good enough for the central purposes of a moral theorist if it should turn out—as I hope it has turned out—that some of the kinds of facts I pointed to earlier are conclusive reasons for *believing* that *A* ought to *V*. Whether they are therefore in some appropriate sense conclusive reasons for *A* to *V* is a question we can leave aside.

12.

The suggestions I have made obviously do not constitute a theory about what a person ought to do. What I have wanted to do is only to set out some features of the structure that I think such a theory should have. It should take account of the multiplicity of the ways of being good and bad, and it should tell us how the ways of being good and bad bear on what a person ought to do. The resulting theory would of course be more complex than Consequentialism, but that is only to be expected.

Filling the structure in requires supplying an account of what rights people have, and what makes it the case that they have them. I have said almost nothing about that here; I have simply helped myself to the notion of a right, in my characterization of justice and injustice in particular.¹³

More generally, it requires supplying an account of what marks an act as giving or failing to give adequate weight to the interests of all affected by it. I have no such account.

On the other hand, I think it unclear what such an account would have to look like—that is, how much precision should be expected of it. Comparisons between gains by one person and losses by another are no-

¹³ My own account of these matters appears in *The Realm of Rights*.

toriously difficult. Extreme examples are easy enough to construct: if Alfred gains relief from a sore throat, he gains something, but if Bertha loses her life, she loses markedly more. How much more can be said than that that difference is *sufficiently great* to mark as miserly an act by Alfred by which he relieves his sore throat with a drug he could easily have given to Bertha, who needs it for life? Again, injustice is breach of a non-overridden right, and the question whether a right is non-overridden turns on how stringent the right is, and how good for others infringing the right would be. Extreme examples are easy enough to construct: Bertha's right to not be killed by another is very stringent, and the fact that Alfred would gain relief from a sore throat by an act by which he would kill Bertha does not override Bertha's right.

There plainly is much room for differences in opinion about what counts as giving adequate weight to the interests of all who would be affected. The kind of theory I see the need of cannot be expected to contain an algorithm for settling such disputes. It has to be allowed to contrast much and little, and to invite us to attend to, and to argue by analogy from, particular cases in which it is clear on any view that a gain by one would be large and a loss to another would be small.

Moreover, it seems to me that such a theory should leave open the possibility that some such differences in opinion are not settleable at all. Not just that there may be nothing that will bring the parties to a dispute to come to agreement with each other; rather, more strongly, that there may be no correct answer to the question under dispute. The fact that there are such disputes on moral matters (if there are) should trouble us no less than does the fact that there are unsetttable disputes on nonmoral matters. No less, but also no more.

If we take practical reasoning to be reasoning about what to do, then conducting it well requires making exactly those more-or-less delicate contrasts between much and little and arguing by analogy from them. Allowing for the possibility of unsettability is allowing for the possibility that in some cases, no conclusion can be shown to be correct. But that leaves plenty of room for cases in which conclusions are reachable—and I see no good reason to think that we are incapable of reaching them when they are.