

The fourth dimension

A. J. Julius, November 9 2009

“Suppose you were prime minister,” says Scanlon,¹ relaying an example from Thomson,

and the commander of the air force described to you a planned air raid that would be expected to destroy a munitions plant and also kill a certain number of civilians, thereby probably undermining public support for the war. If he asked whether you thought this was morally permissible, you would not say, “Well, that depends on what your intentions would be in carrying it out. Would you be intending to kill the civilians, or would their deaths be merely an unintended but foreseeable (albeit beneficial) side effect of the destruction of the plant?” Holding fixed the actual consequences of the raid and what the parties have reason to believe these consequences to be, might an action be permissible if performed by an agent with one intention but impermissible if performed by an agent with a different strategy in mind? I agree with Thomson in finding this implausible.

Are they right?

I

A change of example will help. (My imagination is refusing to play war.) By administering a strong dose of morphine to a patient who’s certain to die within hours you will both relieve her pain and cause her to die sooner, releasing yourself to an early start on your ski vacation. The reason you have to relieve the pain is, I’ll suppose, far stronger than the reason you have to prolong her life in pain for a few hours. It seems to me that

(1) You should give her the morphine.

and so that certainly

(2) It’s false that you should not give her the morphine.

On my way to those conclusions I didn’t stop to consider your beliefs about reasons, and in fact it seems to me that

(3) The truth of (2) is constitutively independent of your beliefs about reasons for giving the drug.²

What room remains for common-sense morality’s condemnation of killing intended to let you leave for your vacation? Consider

(4) If you believe that the fact that giving morphine to a person will kill her and allow to leave for your vacation is a reason to give her the morphine, then your having that belief makes it the case that you should not give her the morphine.

If (4) is true, then (2) can’t be true of a person who has that belief. This denies the independence declared in (3).³

But (4) is not the only possible representation of a moral prohibition of intentional killing. Consider

(5) You should not (give morphine to a person and believe that the fact that giving morphine to her will kill her and allow you to leave for your vacation is a reason to give it to her).

In (5) the *should not* takes as its object, not a single act, but the (parenthetical) conjunction of an act and a belief. And it asserts that *should not* unconditionally on the actor's beliefs. So it's consistent with (2)'s being true of you whatever you believe. It can be at once true of you: that you should not (give the drug and believe that the vacation provides a reason); that you believe that the vacation provides a reason; and that it's not the case that you should not give the drug.

II

The move from (4) to (5) looks, well, like a move. A trick too cheap to solve any problems. Half the time it still looks that way to me. Now and then, and including *now*, I think there's something to it.

I think that, beyond the facts as to whether I should hold various beliefs and as to whether I should perform various acts individuated independently of the beliefs that accompany them, there are facts as to whether I should instantiate various act-and-belief combinations and as to whether my thought and action should follow various act-and-belief sequences. And, while we're at it, facts as to whether a single actor should perform sequences of acts. And, if you don't mind, facts as to whether the members of a population of actors should do various sets of their several acts. There are several *levels* to what we should do and how we should be. Above a ground level populated by single-act and single-attitude requirements, we face requirements to seek or avoid compounds and series of those ground-level things. Reduction is tempting, as usual. Maybe the compound requirements hold in virtue of or consist in conjunctions of the simple ones. Or maybe dependence goes in the other direction; it is rumored that, for every act that I should do, I should do it because it forms part of an extended acting-on-a-maxim that's required of me. But until you accept such a subordination, you have a competition: compound requirements vying with simple ones to regulate the objects that figure in both.

The problem of how I can answer at once to (1) in respect of my acts and to (5) in respect of my acts-and-beliefs is an example. Let me work through it a bit further. To focus this discussion I'll be asking whether (5) succumbs to versions or relatives of the objections that Scanlon raises against (4).

III

A

For an object to be permissible or impermissible, Scanlon claims, it must be that the actor can choose that object. You can't choose which reasons beliefs to have. Nor can you choose to make it the case that those beliefs cause your acts. So you can't choose whether or not to act for a bad reason. And it can't be that acting for a bad reason is impermissible.

If this voluntariness constraint on permissibility is ratified, the *should not* of (5) is not the *should not* of impermissibility.

I'm happy to go along. I am going to suppose with Scanlon that there's a kind of thinking about whether to perform an act that's normally and properly governed by consideration of reasons and principles that take individual acts as their objects. And the fact that an act is impermissible, I'll accept, just is the fact that, according to one of (a true and specifically moral subset of) those single-act principles, it shouldn't be done. So (5), whatever it's about, is not about permissibility.

B

If a doctor who has the vacation reasons belief gives the morphine, she violates (4). If she withholds it, she violates (1).

Maybe you believe in tragedy: it can happen, you think, that every available act violates some moral requirement. But that's easiest to believe where the requirements spring from mutually independent sources—the interests of two different persons, for example. It's harder to accept that no act available to this doctor is permissible in virtue of her single patient's survival and relief. And for this reason it's hard to accept (4) together with (1).

But putting (5) in place of (4) springs the trap. For, again, it can be at once true of a person: that she should not (do x and believe that p); that she believes that p; and that it's not the case that she should not do x. Tragedy averted.

C

Some people balk at the last step. Some simply deny the compatibility of those three things. But a different objection takes us back to the voluntariness-permissibility connection.

I claim that it can be true of you, even as you hold the vacation belief, that your giving the drug is permissible. But if you're to exercise that option without violating (5), it won't be enough to give the drug. You'll also have to throw over your belief. Your only hope of upholding (1) and (5) together has you first changing your mind about the reasons and then giving the drug. Since you can't choose to change your mind, this result has led a couple of people to object that my view violates Scanlon's voluntariness condition on permissibility after all.

On its face this reaction is just mistaken. While I say that the mind-changing, drug-giving sequence is required, I don't say that giving the drug with your mind unchanged is impermissible. At Scanlon's suggestion I reserved permissibility for the act itself, and like him I deny that the act depends for its permissibility on a change of mind. (2) tells you that it's permissible, whatever you believe, and (5) does not disagree.

D

Behind the objection this more interesting thought about voluntariness might be at work:

(6) In every situation there's some act that a person can choose such that by choosing that act she can satisfy all of the moral requirements that apply to her action in that situation.

This takes up the no-tragedy condition while insisting on choice. It insists, not merely that the objects of *permissibility* considerations be choosable, but that all moral requirements on action be meetable by choices. And this could spell trouble;

choosing to give the drug doesn't suffice for (5) while choosing not to give it violates (1).

Not every way of choosing to give the drug, it is true, suffices to meet (5). But some choices of that act will do it. In particular, good ones. If you ask yourself whether to give the drug, and if all goes well—if you consider which things count in favor or against, and if you conclude that relief and only relief counts in favor while survival counts against, and if you judge on the basis of those that you should give the drug, and if that judgment causes you in the right sort of way to intend to give it—then all of that constitutes, I think, your choice of the act. The act and belief state required by (5) and (1) is not something you can bring about by choosing *that state*. But it's something you can bring about as part of your choosing *the act*, which you can choose. And that possibility of satisfying (5) and (1) by choice is enough, in the eyes of (6), to qualify them as genuine requirements on action.

E

Is this really enough? If voluntariness is still troubling you, you might have this stronger constraint in mind:

(7) In every situation, every mental transition that a person must undergo if she is to satisfy all of the moral requirements that apply to her action in that situation is itself a choice.

What's really bothering you, I imagine, is that the change of belief that's necessary to your satisfying (1) and (5), though it may form part of a choice, is not itself a choice. And (7) demands that each of the changes necessary to a person's satisfying the moral requirements on her action be itself a choice.

But (7) is no good. If I don't believe that I should save the drowning child, if I don't believe that anything counts in favor of reaching out to her, then I can't come to choose to save her unless I change my beliefs. That change is not a choice. But save the drowning child I should.

F

A final diagnosis to try out. The requirement (5) that you not have the belief-and-act combination, and the requirement (1) that you do the act, are requirements on action. You might think that, because you can satisfy these requirements only if you are free of the belief, these requirements make it the case that you should not believe the thing. That consequence seems unacceptable, as this seems to say:

(8) If requirements on action make it the case that you should be in some state, then you can choose to be in that state.

But I don't suppose that (1) and (5) make it the case that you should not have the reasons belief. I think that facts can help to make it the case that you should believe that *p* only by bearing on *p*'s being the case.

G

I'm still not sure which point Scanlon wants voluntariness to make. I apologize for trying on so many possibilities so tediously. My conclusion is that none of them fits (5).

IV

A

Voluntariness is not the only thing bothering Scanlon about (4). Another misgiving seems to involve *objectivity*.

The charge against (4) might just be that my having an attitude can't make a normative requirement true of me. But that way of insisting on objectivity is not sane. The fact that I love my mother, which makes my presence valuable to her, can help to make it the case that I should stay home rather than join the resistance. Our requirements are not generally independent of our attitudes. The dependence on attitudes that's supposed to underwrite (4) is suspect, if at all, only in some special way beyond its being such a dependence.

B

Dialogue's distracting. Consider a Thomsonian soliloquy:

Should I give the morphine? That depends on what I'd be taking as a reason for giving it. I should do it if I'd be taking pain relief as a reason.

But I shouldn't do it if I'd be taking an early vacation as a reason. So what would I be taking as a reason?

What's weird about the speech is that the speaker takes the answer to her deliberative question—should she do it?—to depend on her having or lacking attitudes—among others, her beliefs about reasons—that she's forming or reforming within the same course of thought through which she's trying to answer that question. Even the prince of Denmark was never so self-involved. Normal reasoning about what to do seems instead to take as its object normative facts that you take to hold independently of your performing the reasoning and of its results in your intentions and beliefs. The not so normal exceptions are important; I'll get to them soon. But this is how I experience a certain kind of normal thinking about whether to do a thing.

(5) is true to that experience. Unlike (4), (5) leaves open that your deliberation about which reasons count in favor of or against giving the morphine and about whether you should give it amount to the discovery of facts all of which hold independently of your deliberation and of the attitudes it produces or refines.

C

A principle that tries to formulate this kind of objectivity:

(9) The fact that you should x is constitutively independent of your beliefs about whether you should x and of your beliefs about such matters as help to determine whether you should x.

But this is still too much. My belief that the cause is just, because it makes me an effective combatant, can help to make it the case that I should join the resistance—a fact that also depends constitutively on the *fact* that the cause is just. But (9) would insulate the fact as to whether I should join even from this faultless dependence on my belief. It's not always a mistake to invoke, at some late stage in your reasoning about whether to do something, the fact that earlier in that reasoning you'd reached some conclusion about reasons for it.

D

In that example my reasoning about whether to fight breaks into my settling of three issues: whether the cause is just, whether I'd be good at advancing it, and whether I should join a just cause that I'd be good at advancing. The determinants of these issues are mutually independent; what makes the cause just, or what makes its justice and my high fighting capacity sufficient for my joining it, involves nothing of what makes me a good combatant. And vice versa. Their independence, I imagine, makes it okay for me to regard my belief about the first issue as fixed—as invariant with respect to the thinking about the second and third issues that lies ahead of me. The fact that I regard the second and third issues as open sits comfortably with my treating the first as settled and so with my placing my position on it out of the reach of my current reasoning.

This test is failed in other, objectionable instances of attitude dependence. Suppose I were to take the fact that I want to climb the mountain as helping to make it the case that I should climb the mountain. Then I'd be taking as fixed an attitude, my wanting to climb, that depends for its appropriateness on the very matters that determine whether I should climb: the danger, the view, my responsibilities to people on the ground. I shouldn't regard the question of whether I should climb as open without regarding those issues as open and so without treating as open the appropriateness of my wanting to climb. And so it'd be a mistake to treat that attitude as invariant with respect to my further deliberation about whether to climb.

Notice that no such problem arises where I take the fact that *you* want me to climb as helping to make it the case that I should climb. That attitude of yours *is* invariant with respect to my reasoning about whether to climb, and I make no mistake in treating it that way. This helps to explain the otherwise puzzling appearance that, between my wanting to climb and your wanting me to climb—between these two attitudes in favor of the same (ahem) proposition—only one can help to determine whether I should climb.

E

Back to the doctor. Consider two deliberative issues: whether the ski weekend counts in favor of giving the drug and whether you should give the drug given that you believe that skiing is a reason. These issues are not independent in the sense just set out. Whether you should give the drug given that you take skiing as a reason depends on whether it is a reason. If you think both that skiing is a reason for drug giving and that your having this belief counts against your giving the drug, then you have some more thinking to do: you ought to reopen the whole question of whether skiing counts in favor of drug giving. It's a mistake to press ahead with your deliberation about the second issue while taking as a constant of that deliberation the fact that you've answered the first as you have.

F

Suppose that taking your belief as fixed in this way is indeed a mistake. I want to conclude that your having the belief does not help to determine whether it's the case that you should give the drug. It does not count against your giving it. For if some fact helps to make it the case that you shouldn't do some act, it can't be that your counting it as a reason against the act is a mistake. The circumstances that set the boundaries of our practical problems and that ground the requirements that

apply to them can't include the fact that I have some attitude unless I can regard that fact as constant and still deliberate well.

I propose that (4) goes wrong not by violating objectivity-as-independence-from-attitudes but by putting reasoners in the bad position of treating as fixed attitudes that they ought not to take as fixed—by trespassing against a principle like

(10) If the appropriateness of holding an attitude A at t depends on matters that you should be attempting to settle in your thinking, at t and beyond, about whether to x at t + 1, then the fact that you have A at t can't help to make it the case at t that you should x at t + 1.

For its part (5) passes this test that by eschewing any recommendation of acts conditional on beliefs.

G

Let's train this proposal on another of Scanlon's examples. You buy some poison. If you're intending to use it to kill your lawyer, then this intention makes it the case that your act of buying poison will probably contribute causally to your lawyer's death. Scanlon concludes that, because it's impermissible to knowingly help cause the lawyer's death, buying the poison is impermissible in virtue of that intention. If you intend just to kill weeds instead, buying it is permissible since that purchase probably won't contribute to any person's death. Scanlon offers this reasoning as a limited vindication of motive-dependent impermissibility. The "predictive significance" of your murderous intent for a calculation of the act's consequences presents a special and *derivative* respect in which this attitude can matter to the act's permissibility.

H

The test on the table suggests that, in conceding this, Scanlon concedes too much. A person who intends to kill even as she takes the fact that buying poison contributes to a killing as a reason not to buy has, again, some more thinking to do. You should reconsider both positions. It's a mistake to plow ahead in your reasoning about how to avoid contributing to a death while taking as given your intention to cause one. And so your having that intention can't help to make it the case that you shouldn't buy the poison.

By turning from simple to compound principles, we can salvage a deliberative verdict like Scanlon's while respecting (10).

(11) You should not (buy poison and intend to use it to kill your lawyer) which in turn holds in virtue of

(12) You should not (buy poison and use it to kill your lawyer).

To derive, from (11) and from the fact that you intend to kill, a requirement that you not buy poison is no good. This would have you holding as fixed an attitude that you ought to suspend in flux. But if you start by accepting (12), the principle that grounds (11), you won't come to buy poison as part of a plan for the sequence of acts that violate (11). You won't come to buy with the intent to kill.

V

A

The deliberation by which I come to believe that the vacation is or isn't a reason is just some thinking about whether that's the case. Does it seem to be true? What are some grounds for it? And then, as I've conceded to Scanlon, when I'm deciding whether to give the morphine, I'm considering the pain and survival reasons and not (5) itself. My formation of the reasons belief, and my formation of the intention to give the morphine, are properly the subjects of their own deliberative processes, no step of which has me considering or invoking (5).

The complex principle (5) seems in this way *crowded out* of my deliberation about the individual acts and attitudes by single-act and single-attitude requirements. In a typical piece of the reasoning through I come to comply, the principle (5) does not occur as the content of any step.

This observation provokes another doubt about (5). It casts (5) at odds with
(13) In any situation in which you should x, you can come to x as the outcome of good reasoning one of whose steps includes the thought that you should x.

This is an ought-implies-can-be-guided principle. It requires that normative requirements be guiding in a very strict sense: it insists that a person's recognition of any such requirement can help, as a step in some good reasoning of hers, to bring it about that she fulfills it.

B

The principle is obviously false. You should intend to feed the hungry traveller. But you can't come to intend that as the result of good reasoning that invokes the fact that you should intend it. You should believe that it's sunny in Los Angeles. But you can't come to believe that as the result of good reasoning that invokes the fact that you should believe it.

Perhaps (13) is true only when its x ranges only over acts. Normativity comes in different flavors, for acts on one hand and for intentions, beliefs, emotions on the other. Only the first normativity and not the second is made up of requirements that guide in the strict sense of (13).

It's unclear where the suggested restriction would leave us on (5). (5) seems subject to (13) insofar as it purports to govern an act. But it seems to duck (13) insofar as it purports to govern a belief.

In any case I doubt that we should settle for this relegation of (13) to the side of acts alone. The hypothesis that act- or intention-governing practical shoulds, and belief-governing epistemic shoulds share a common nature as the normative relations in which we stand to our various doings and beings is too powerful to be discarded before we've made at least some effort toward a generalization of (13) that covers the normative waterfront.

C

Where you should intend to feed the traveller, this is in virtue of the fact that you should feed her. The second fact is a ground of the first. The remarkable thing is

that, while your belief that you should intend to feed her can't lead to your intending that, your belief that you should *feed her* will get you there. A contemplation of the *ground* of this requirement is a step in some good reasoning through which you can come to fulfill it.

Where you should believe that *p*, this is in virtue of certain facts: the fact that *p*; or facts that count as evidence in favor of *p*'s being true. If you have a thought, or a perceptual experience, that takes as its content the fact that *p* or the facts that are evidence for *p*, you either already comply with the requirement that you believe that *p*, or you can reach compliance with it by a good mental transition from these thoughts or experiences.

D

In face of the normativity of intention and belief, we can't expect every derived normative fact to guide us in every situation to which it applies by itself forming the content of a step in the process by which we come to comply with it. For that expectation I propose to substitute this weaker requirement of guidingness:

(14) In every situation in which you should *x*, you can come to *x* as the outcome of good reasoning one of whose steps takes as its content one or more of the grounds of the fact that you should *x*.

This ensures that the set of all normatively relevant facts—all of the facts of what you should do and how you should be and all of the facts that ground those ones—is capable of guiding you toward compliance with every requirement in that set. Every one of those requirements can be fulfilled as the result of reasoning one of whose steps take as their contents members of the set.

E

This is a big and sketchy proposal. I won't say anything more for it now. Let me just draw the moral for (5).

If (5) is true, it's true partly in virtue of the fact that an early start for your vacation is no reason for giving the drug and of the fact that giving the drug is opposed by the reason of survival. If you invoke these facts to yourself, and if you consider the fact that pain relief counts in favor of the drug, and if you conclude on the basis of these that you should give the drug, then you'll have come to give the drug without holding the vacation belief. These grounds of (5) can figure as the contents of thoughts that form a chain of reasoning at whose end you satisfy the principle. So (5) satisfies (14) as well.

VI

I've offered some ways of understanding the apparent problems of voluntariness, of objectivity, and of guidingness on which they turn out to make no real trouble for (5). In deciding what to make of these proposals, you might also consider that (5) is not alone in raising these problems by governing acts together with attitudes.

A

Another Thomson story. Because of a freak electrical disturbance, flipping a light switch in your apartment will cause a lightning flash that kills your neighbor. You

don't know this. Your earlier experience with the switch makes you justified, let's assume, in believing that turning it on is harmless.

Thomson claims that it's impermissible to flip the switch. This is presumably in virtue of something like

(15) If flipping the switch will hurt someone, you should not flip it

Scanlon denies this. He defends instead a principle like

(16) If the possible belief that flipping the switch will be harmless is not justified in your evidentiary situation, then you should not flip it.

and concludes that flipping the switch is permissible after all.

In defense of (15) Thomson imagines another dialogue. Suppose B has run the test and knows that switching is harmful.

A: Should I turn on the switch?

B: If I tell you what I know, then it'll be true that you shouldn't. But so long as I don't, it's false that you shouldn't.

Scanlon says that what's wrong with B's speech is that it presupposes that A's coming to believe that the act is harmless will change its permissibility. And of course that's not true on Scanlon's (16) since permissibility depends not on whether A believes that the act is harmless but on the justifiedness of that possible belief.

B

But B's telling you that switching is harmful won't merely change what you believe. By changing your evidence it will also change the justifiedness of your belief that it's harmless. This conversation makes that belief unjustified and so by (16) it makes turning the switch impermissible. So what B says is true after all on Scanlon's view. But it still seems mistaken.

Suppose that, despite being justified in your belief that switching is harmless, you run the test anyway and find out that it's harmful. On Scanlon's view, by running the test, you make the act impermissible. But surely that's wrong. Surely you *discover* that it's impermissible.

There's evidence that Scanlon agrees. He says that when a person knows that switching is harmful "it is the fact that is known, not the fact of his knowing it, that he is to take as counting against his action." But this diagnosis seems excluded by Scanlon's (16). On (16) what counts against the act is neither the fact of harmfulness nor the fact that you believe it harmful but rather the fact that a possible belief in harmlessness is not justified. So in this last passage Scanlon depicts the person as acting on the basis of considerations apart from the ones that his own principle makes determining of permissibility.

C

I have so far only an ad hominem point against Scanlon. By endorsing (16) he endorses a principle that seems crowded out of a typical piece of deliberation about whether to flip the switch on his own characterization of how that deliberation goes.

But Scanlon is surely *right* that Thomson's (15) is not the whole normative story. We seem to face another requirement. "One ought not to do what one sees, or should see, will cause serious harm to someone. One ought to take due care not to cause harm." This can be seen, Scanlon says, as follows:

If any principle that it's not reasonable to reject would require us to take C as a consideration counting strongly against an action, then, since it may not be immediately obvious whether C obtains or not, in considering any other principles it would not be reasonable to reject, [we should conclude that one of those other principles is that] we are required to be on the lookout for C, and to take reasonable steps to find out whether or not it obtains. [AJJ's interpolation]

I think the further principle might be this:

(17) You should not (flip the switch and fail to believe with justification that flipping the switch is harmless.)

D

Like (16), (17) is often crowded out. If you start out to think about whether flipping the switch is harmful, you'll be guided not by (17) but by the evidence for and against harm and by the requirement that you form beliefs on the basis of evidence. And if you then conclude that flipping is harmful, you will use (15), not (17), to decide you against flipping. This is some independent evidence that a principle can be genuinely normative despite being crowded out.⁵

E

Suppose on the other hand that your evidence warrants neither the conclusion that switching is harmful nor the conclusion that it's not. Then (15) fails to tell you what to do as you don't know which of your acts upholds it. But (17) is still in force. It offers two options: you can go on trying to find out whether switching is harmful so that, if you conclude that it's not, you can switch consistently with (17). Or you can simply not switch.

In the second case you'd be taking the fact that you don't have a justified belief in harmlessness to count in favor of not switching. Does this violate the test of section IV? No. For again you can decompose your question into two issues. Is switching harmful? And is switching supported or opposed by any facts about what you believe about harm? In taking the second question to be open you need not hold the first open.

F

Notice that on this picture (15) and (17) are consistent even if "should" is understood univocally in them. There is no need to distinguish, for example, a subjective from an objective use of "should".

But if Scanlon's right about permissibility—if an act's being permissible consists in the fact that, according to the single-act principles that you should use in deciding whether to do it, it's false that you shouldn't do it—then we should take (15) and not (17) as owning permissibility.⁴

VII

A

What kind of moral fact, then, is the fact you should not (switch and fail to believe with justification that switching is harmless). Scanlon identifies three moral dimensions: permissibility, blame, and meaning (the significance for your relations with others of the way in which you decide how to treat them). (17) is not about permissibility, since it's not the principle you're to consider in deciding whether or not to do the act. It's not about blame since it makes a normative claim that you shouldn't be in a particular state. And I don't think it's only about meaning. A person who flicks the switch despite her belief that that's harmful does not merely reason incorrectly about whether to flip the switch; she also mistreats her neighbor by failing to "take due care."

I think we need another dimension, then. I think that it's probably *wrongness*. That's going to be a hard case to make. Here I'll just point out that Scanlon's own contractualist theory of wrongness might be recruited to make it.

B

For Scanlon an act's being wrong is its violating a principle that no one can reasonably reject given that she shares with others the aim of finding principles that no one can reasonably reject given that shares with others ... An act can violate a principle only if the principle takes a simple act as its object. Principles like that are, as I've allowed, the ones that establish permissibility. So it's credible, on this picture, that an act is wrong if and only if it's impermissible.

But are acts the only things that can be wrong? It's not clear that Scanlon is in a position to insist on that.

C

Scanlon offers his theory as explaining the reason-giving force of the consideration that an act is wrong. In some places he explains that by appealing to a basic reason that each person has to avoid acting in ways to which other people can object with respect to reasons based in their own interests and claims. I've criticized this explanation in another paper.

In other places he offers what I take to be a better one. He argues that there's positive value in the fact of several persons' doing things to and with one another in ways that are governed by each person's recognition of the reasons that each is given by the others' interests and claims. The fact, as Scanlon puts it, of mutual recognition. You should avoid doing a wrong act, on this view, exactly because doing it is incompatible with mutual recognition.

D

But here's something else that's incompatible with that: interacting with people in a way that's not appropriately governed by consideration of the reasons arising from their lives. In particular, performing a permissible act as the product of reasoning that takes bad account of those reasons. (Notice that the reasoning itself, or my holding false beliefs about reasons, is not incompatible with that. They don't by themselves make it the case that I'm not interacting with others in ways governed by the right sort of reasons. For I might yet reason my way out of them

and act for the right reasons. Once I've acted on the wrong reasons, on the other hand, it's too late for mutual recognition.)

E

Scanlon's very wise strategy for finding out what wrongness is, set out early in *What we owe to each other*, is to identify wrongness with whatever property turns out to best explain our having decisive reason not to perform the acts that we judge to be wrong. Let's run with that. On the version of his explanation that I've just sketched, that property is incompatibility-with-mutual-recognition. And I've claimed that acts are not the only objects that can have this. Act-and-belief complexes in virtue of which I'm treating you in a way that's not governed by the right sort of consideration of your interests and claims are another bearer of it.

If that's right, they're wrong. Welcome to the fourth dimension.⁵

Notes

1. *Moral dimensions*, 2008. In the book's first two chapters Scanlon makes a three-fold case for permissibility's independence from things like motive or intent. For one thing he gives the general deliberative arguments for it that I consider in the text. For another he reports that he can't see any good explanation of permissibility's depending on intentionalness or reasons beliefs in the central cases for which double-effect-style principles are taken to condemn intentional killing. For a third thing he has an error theory that's supposed explain away our temptation to think that permissibility depends on intent.

The error theory in turn has two parts. The first draws on the following picture of moral principles.

Principles are facts about which considerations count in favor of which acts and facts about which acts you should do that hold in virtue of those reason facts. You can use these deliberatively—use them to make up your mind whether to do the act—by considering the reasons and requirements in a way that leads you to form the attitudes and perform the acts that they recommend. Or you can use them critically—use them to assess a person's reasoning—by asking whether the person properly takes the principles into account in the reasoning that leads to her doing what she does.

Scanlon claims that people are sometimes confusing the deliberative and critical uses. Thinking critically they've concluded that the commander who bombs for demoralization has failed to take the right considerations into account. It's easy to slide from this into the thought that she acts impermissibly. But that's a mistake—a slide from critical scrutiny of the commander's reasoning into a deliberative concern with what the commander should do.

The other part explains a few instances of motive dependence as “special cases”. See for example the poison example discussed in IV.G.

2. Like Scanlon I'll suppose that a central way in which some aspect C of an act x can count as intentional is for the agent to believe that C provides a reason for x'ing and for that belief to help cause her x'ing in a normal sort of way. On that

assumption (3) asserts the independence of a fact about whether to do something that causes a death from the intentionalness of killing. If you disagree, then forget about intentionalness and focus instead on this question: Can an actor's believing something to be a reason for her act help to make it the case that you shouldn't do the act?

3. You might object to (4) on the independent ground that giving the drug is wrong only if the reasons belief about skiing is *decisive* in the doctor's decision to give it. If that bothers you, feel free to weaken the principle accordingly. For example:

(4') If it's the case that, if you were to come to intend to give the drug, your belief that a ski vacation is a reason to give the drug would contribute decisively to your coming to intend to give it, then you should not give it.

Because the objections to (4) that I'll be considering also stand against (4'), and because my arguments that (5) avoids those objections will extend to an corresponding weakening of (5) like

(5') You should not give the drug as the outcome of deliberation in which the belief that the ski vacation is a reason for giving it is decisive.

I think it OK to focus on (4) and (5).

4. But maybe Thomson's example makes things too easy? Here's a harder one.

Ten miners are trapped in a single shaft, in A or in B. You don't know in which, and you're justified in assigning equal probability to each alternative. If you divert the flood into one of the shafts, it kills any miners who are in there. If you do nothing, the water fills both shafts more shallowly and kills only one miner.

We want to say that, knowing what you know, it is true of you that

(18) You should flood neither shaft.

but also that

(19) If the miners are in A, you should flood B.

But if in fact

(20) The miners are in A.

you seem to have

(21) You should flood B.

Philosophers have tried to win consistency by assigning an equivocal interpretation the shoulds of (18) and (19), by denying the objective (19), or by denying the subjective (18) outright. The discussion of Thomson suggests an alternative: maintain (19) while revising (18) in favor of

(22) You should not (flood B and fail to assign, with justification, a credence less than alpha to the miners' being in B.)

This lets "should" function invariantly. It makes sense of your consulting (19) as a way of finding out whether to flood the shaft. But it also implies that you violate a moral requirement—(22)—when you flood the shaft while remaining uncertain that it's empty.

Given your information, you are condemned either to violate (21) or to violate (22). Is that an unacceptable result? No. The two principles are practically consistent in this sense: There's no state of information such that you'll have reason to accept (21) and such that you violate (22) if you flood shaft B. So there's no state in which they offer conflicting recommendations both of which you'll hear.

5. Thanks to David Bordeaux, Lindsey Chambers, Jakob Elster, Mark Greenberg, Kevin Falvey, Matt Hanser, Barbara Herman, Pamela Hieronymi, Mike Martin, Thi Nguyen, Michael Rescorla, Ira Richardson, Tim Scanlon, Wally Siewert, Gary Watson, Stephen “UCLA” White, and Andrew Williams.