

You Gotta Do What You Gotta Do

JOHN GIBBONS

University of Nebraska—Lincoln

There are two separable issues about the role of the mental in the determination of practical reason. We can ask about the role of desire, and we can ask about the role of belief. Beginning with the latter issue, we can ask whether it's the belief that you're out of milk or the fact that you're out of milk that gives you a reason to go to the store. And we can ask whether it's the belief or the fact that makes it reasonable for you to go to the store. Just to lay my cards on the table, I think that good reasons are things that make things reasonable. So we cannot give an account according to which the reasons are determined by the facts while what's reasonable is determined by the beliefs. We need a unified account of reasons and rationality. And the best unified account is not in terms of the facts. It's in terms of your beliefs, or your psychological states, or elements of your perspective. We will return to this issue, but for the most part, it will remain in the background.

Suppose you know that you're out of milk, so you have both the belief and the fact. You still don't have enough to get you to the store. What do you need to add? A natural first suggestion is that you need to add a desire: a desire for some milk. If you know that you're out of milk but don't want any more, you might not go to the store, or have any reason to go to the store, and you might reasonably refrain from going to the store. Generalizing from this case, we might say that desires play the following three roles: they're necessary for causing intentional action; they're necessary for providing reasons to act; and they're necessary to make actions reasonable. I haven't said much to motivate this view. But that doesn't mean that there's nothing to say. At this point, we're just looking at options.

And options are not hard to come by. You might think that beliefs sometimes play one or the other of the roles reserved for desire by our first option.

Or you might think that beliefs sometimes play all of these roles. You might think that beliefs always or almost always play some or all of these roles. Once you've chosen a role and a frequency with which beliefs play it, you still have your choice of beliefs. Maybe the normative belief that you ought to get milk could play the relevant role. Maybe the evaluative belief that getting milk would be good could play the relevant role. And maybe the psychological belief that you will want milk could play the role. Finally, in addition to the normative, evaluative, and psychological beliefs, there are the corresponding normative, evaluative, and psychological facts which might sometimes or always play some or all of the roles.

In order to make some progress, I will focus on one relatively small corner of this landscape. Can the normative belief that you ought to ϕ ever, without the help of desire, rationalize or make reasonable a desire or intention to ϕ or an attempt at or intentional action of ϕ -ing? After some attempt to clarify the issue, I'll look at an argument for a negative answer to our question, present a couple of arguments for an affirmative answer, and look at some of the consequences of an affirmative answer for our wider landscape.

The Humean Theory of Motivation

According to Michael Smith, the Humean Theory of Motivation is a claim about motivating reasons, not normative reasons.¹ Informally, the idea is that all motivation has its source in the presence of a desire and means-end belief. But the idea that Smith thinks is crucial to the Humean theory is put more formally as follows.

(P1) R at t constitutes a motivating reason of agent A to ϕ iff there is some Ψ such that R at t consists of an appropriately related desire of A to Ψ and a belief that were she to ϕ she would Ψ .²

According to Smith, motivating reasons consist of beliefs and desires, but that doesn't mean that normative reasons do so as well. I'm not sure I understand the distinction between motivating and normative reasons. Smith tells the following story to illustrate the distinction.³ The painting in front of Smith is a Picasso, but he doesn't know this, and there's nothing in the story to suggest that he should have known. Smith wants to buy a Picasso. So, according to Smith, he has a normative reason to buy the painting; there is a "requirement of rationality" that he buy the painting; and he is "rationally required" to buy it.⁴ Given how much Picassos actually cost, spending that much money on a painting you have no interest in is clearly and obviously irrational. Whatever else you're rationally required to do, you're rationally required to be rational and rationally prohibited from doing clearly and obviously irrational things. So whatever this normative reason is supposed to be, it doesn't look like a requirement of rationality.

Smith tells us that there is something that both motivating and normative reasons have in common in virtue of which they're both reasons: citing either renders an agent's action intelligible.⁵ So suppose Smith buys the painting not knowing it's a Picasso and not liking it for any other reason. The unknown fact that it is a Picasso can no more make the action intelligible than it can make the action reasonable. Normative reasons, according to Smith, are facts about what would be desirable for you to do.⁶ Right now there are many unreasonable things that you could do that will, in the fullness of time, turn out to have been desirable. You have no way of knowing what any of these things are. If you knew they would be desirable, it wouldn't be unreasonable for you to do them. But given what we know about causation and how far into the future the consequences of our actions can reach, we can be reasonably sure that many of the stupid things people do will eventually turn out to be for the best.

Normative reasons or facts about what will eventually turn out to be desirable are completely independent of the rationality and intelligibility of action. So what about motivating reasons? Do they make things reasonable or determine the rationality and intelligibility of action? It certainly seems that way. Suppose that Smith does not buy the painting he has no interest in. It's not just that this is a reasonable thing to do. It's that this is reasonable because he has no interest in the painting. It's made reasonable and intelligible not by the facts but by his take on the facts, or his perspective, point of view, or whatever you want to call it.

I think that this is a perfectly natural way to think about motivating reasons. Suppose you desire E and think that doing M will get you E. What is it about this belief-desire pair that makes it a reason to M? It seems that the pair counts as a reason, a motivating reason, in virtue of the fact that it makes doing M reasonable or intelligible from the agent's point of view.⁷ Of course, as Davidson says, there may only be a weak sense in which the belief-desire pair makes the action reasonable. If either the belief or the desire is unreasonable then doing M will be unreasonable as well. But whatever else is wrong with M, the problem is not with the transition from the belief-desire pair to the action. So to say that the motivating reasons make the action reasonable in the weak sense is to say that the transition from them to the action is a reasonable move for the mind to make. To say that the motivating reasons determine the rationality of action is to say that if the relevant mental states are reasonable then so is the action, and if they're unreasonable, then so is it.

You get the same thing in the theoretical case. Suppose you believe that p and that if p then q. These could lead you to believe that q. Whatever else is wrong with believing that q, the problem is not in the transition. If, on the other hand, you believed that p and that if q then p, and these led you to believe that q, then the problem could be in the transition, and believing that q could be unreasonable even if the antecedent beliefs were just fine.

But if both the antecedent beliefs and the transition are reasonable then the conclusion is reasonable as well, and it's reasonable in virtue of where it came from.

So there is one perfectly natural picture of motivating reasons that works in both the theoretical and practical cases. Motivating reasons are the kinds of things that make you do it or make you believe it, and they're the kinds of things that determine the rationality and intelligibility of action and belief. The belief-desire pair counts as a motivating reason on this picture in virtue of the rationality of the transition from them to the action. If it turns out that the transition from a set of beliefs (without any desires) to an action is equally reasonable, then the set of beliefs will also count as motivating reasons and the Humean Theory of Motivation will be false. So on this picture of motivating reasons, the Humean theory is not a contingent, empirical claim about what causes action in humans. It's a claim about which mental state transitions are reasonable, something far more appropriate for philosophers to argue about.

Though I think that this picture of motivating reasons and the interpretation of the Humean theory it delivers are both perfectly natural, I'm not at all sure that either the picture or the interpretation are Smith's.⁸ Thomas Nagel imagines a case in which you believe that you'll want something tomorrow even though you don't want it today.⁹ He asks, in effect, whether this belief could make it reasonable for you to prepare now for what you think you'll want tomorrow. It certainly seems that this belief could make the action reasonable, and it certainly seems that the transition from the belief to the action is a reasonable move for the mind to make. If this is enough for the set of beliefs to constitute a motivating reason, the example certainly seems to show that the Humean theory is false.

Smith says that the example does not cause trouble for the theory.¹⁰ You do have a reason to prepare now for what you think you'll want tomorrow, but that reason is a normative reason, not a motivating reason. You should not be misled by the terminology into thinking that just because the belief makes the action reasonable the belief is a normative reason. Normative reasons are simply not in the business of making things reasonable. Normative reasons are supposed to be facts about what would be desirable. A belief about what you'll want tomorrow, no matter how justified, can always be a mistake. So while the belief might make the action reasonable, it can't make preparing for something you don't want now and won't want then desirable, at least not in the sense in which it's desirable to spend vast amounts of money on a painting you have no interest in.

I don't understand why Smith thinks this reason is a normative reason. Presumably, he thinks that it is not a motivating reason because he thinks that a transition from a set of beliefs to an action would never be a reasonable move for the mind to make. So suppose the beliefs about what you will want cause you to go to the store. Smith might say that your going to the store

was unintelligible, that we wouldn't know what you were doing, or that we couldn't understand your action in terms of the pursuit of a goal: getting some milk. Or Smith might say that your getting milk in these circumstances is unreasonable or irrational. The irrationality would not stem from any problem with the beliefs. He could agree that you'll want the milk tomorrow, that you won't want to get it in the morning, and that you're not doing anything better at the moment. The problem would be with the transition from these beliefs to the action. Going to the store in these circumstances would be the practical analogue of affirming the consequent.

If Smith says that going to the store in these circumstances is not intentional because it's unintelligible or irrational, then we simply disagree about which mental state transitions are reasonable. If he says it's not intentional because it wasn't caused by a desire, then the Humean theory is the claim that reasons that involve desires involve desires. This theory is no doubt true. But it's not the insight it purports to be.

Direction of Fit

The Humean Theory of Motivation is the idea that motivating reasons are always constituted, at least in part, by desires. Here's how Smith argues for the view.¹¹

- (a) Having a motivating reason is, *inter alia*, having a goal.
 - (b) Having a goal is being in a state with which the world must fit.
 - (c) Being in a state with which the world must fit is desiring.
- (HTM) Having a motivating reason is, *inter alia*, having a desire.

The argument depends on the idea of a state with which the world must fit. If we took this idea literally, it sounds like the idea that the world is under an obligation to bend itself not only to our will, but to our every whim, even when those whims conflict. But we're not intended to take this idea literally. The idea that desires are states with which the world must fit is supposed to be an expression of the idea that desires have a certain direction of fit. And the difference between directions of fit is spelled out in counterfactual terms.

Very roughly, and simplifying somewhat, [the difference between beliefs and desires] amounts, *inter alia*, to a difference in the counterfactual dependence of a belief that p and a desire that p on a perception with the content that not p: a belief that p tends to go out of existence in the presence of a perception with the content that not p, whereas a desire that p tends to endure, disposing the subject in that state to bring it about that p.¹²

This is intended as an account of what differences in direction of fit come to. It's only intended as a rough, simplified account. So it shouldn't be too

surprising if we can find some cases where it doesn't quite work. So suppose that I hope that my team won the game last night. So I have a pro-attitude toward the proposition that my team won the game last night. Now suppose I perceive that not-*p*: I hear on the radio that my team didn't win. Is my hope supposed to persist and dispose me to make it the case that my team won the game last night? It's not just that hopes don't always do this. It doesn't look like they ever do this. When I find out that my team lost, I no longer hope that they won.

So there are some cases where this doesn't work. That doesn't show that it's not along the right lines or roughly correct. But if it is roughly correct, we ought to be able to find some cases where it does work. This is more difficult than it sounds. Suppose that you want to eat soon. So you have a desire or pro-attitude toward the proposition that you will be eating soon. Now suppose that you perceive that not-*p*: you find out that you will not be eating soon. What would it be like for the desire to persist and dispose you to make it the case that *p* when you know or believe that you will not succeed?

So maybe it's a matter of fine-tuning the details. Maybe the desire that *p* plus the perception that not-*p* don't dispose you to make it the case that *p*. They dispose you to make it the case that *q*, where *q* is, as it were, the next best thing compatible with the evidence. So you want to eat soon, but you find out that you won't, and this disposes you to make it the case that you will be eating fairly soon, or some time today, or something like that. But we're looking for the crucial difference between beliefs and desires. Suppose I believe that she's home, but find out that she's not. These lead quite naturally to the belief that she'll be home soon. And if I find out that that's false, I'll think that she'll be home fairly soon, or some time today, or something like that. Just as the desire that *p* plus the perception that not-*p* can lead to a desire that *q*, the belief that *p* plus a perception that not-*p* can lead to a belief that *q*, where in both cases *q* is, as it were, the next best thing compatible with the evidence.

Fine. But there's still this difference between beliefs and desires: desires dispose you to act and beliefs don't. Could this be the kernel of truth buried under the counterfactual account? This might be what's behind the counterfactual account. But it's simply a statement of the Humean theory, and as such, it can't be used in an argument for the theory. If this really is a kernel of truth, we need some reason to believe that it's true. It might not take much to convince us that desires dispose us to act. But in the present context, we need some reason to think that beliefs don't and probably that they can't.

So maybe someone can fix the counterfactual account. Or maybe counterfactuals are not the place to look for an account of direction of fit. But surely there must be something to Anscombe's idea that beliefs are supposed to fit the world while the world is supposed to fit our desires. Can't we simply rely on this idea in arguing for the Humean theory without having to worry

about the details of the account? If we're relying on the idea that the world is supposed to fit our desires, and if this idea is supposed to come from Anscombe, maybe we should look at what she says. Anscombe tells a story about a man who goes to the store with a shopping list.¹³ A detective follows him and makes a report of what the man buys. If the man buys what he's supposed to, and the detective's report is accurate, the difference between the shopping list and the report will not be a difference in what they say. As we would put it now, it's not a difference in their content. So what is the difference? Here's Anscombe's answer.

It is precisely this: if the list and the things the man actually buys do not agree, and if this and this alone constitutes a *mistake*, then the mistake is not in the list but in the man's performance . . . whereas if the detective's record and what the man actually buys do not agree, then the mistake is in the record.¹⁴

There's a great deal of room between this and the idea that beliefs are supposed to fit the world while the world is supposed to fit our desires.

First of all, there's an asymmetry in Anscombe's answer. If there's a discrepancy between the report and the facts, it follows that there's a mistake in the report. But if there's a difference between the shopping list and what the man buys, it doesn't yet follow that the mistake is in the action. If you go around Oxford with a list that includes "tackle for catching sharks,"¹⁵ then the mistake is in the list, especially if the list-maker should have known that you can't get that sort of thing around there. If you don't get what can't be gotten, the mistake is not in the action.

If you leave open the possibility of mistakes in the intention, then you must give up the idea that world should always fit our desires. And it does seem that Anscombe is right about this. Sometimes when it seems that all hope may be lost, the thing to do is to persevere, storm the ramparts, and try to change the world. But at least occasionally, when the object of desire is well and truly out of reach, the thing to do is to stop moping, get out of the house, and do all those things that you know will eventually lead to a change in desire. There might be nothing perfectly general to say about when you ought to try to change the world and when you ought to try to change your mind. Perhaps this is why Anscombe never tried to say anything perfectly general. But if there is something to say, it's not that you should always try to satisfy all of your desires, even the conflicting ones, no matter what.

Secondly, Anscombe's story involves an intention (or shopping list), an action, and a belief (or report). It makes perfect, non-metaphorical sense to say that sometimes the mistake is in the action, and sometimes the mistake is in the belief. But when we generalize this and look at cases where the belief and desire are not about actions, we run into trouble. If I want it to rain and it doesn't rain, there's absolutely no reason to suppose that there must be some metaphorical sense to the idea that the mistake is in the weather,

if only we could figure out what that metaphorical sense might be. The project of generalizing from Anscombe's story in an attempt to provide an illuminating account of the difference between beliefs and intentions is a perfectly legitimate and difficult task.¹⁶ But if you think of the task in terms of trying to make sense of the idea that beliefs are supposed to fit the world while the world is supposed to fit our desires, you're not only starting off on the wrong foot. You're not really generalizing from Anscombe's story.

The Quick and Easy Argument

So we haven't yet seen good reason to believe the Humean theory. Is there any reason to think that it's false? As I understand the theory, it's the idea that a transition from a set of beliefs, no matter how extensive, to a desire, intention, or attempt to ϕ , would never be a rational transition. In order to rationally acquire a pro-attitude from a set of mental states, there must be some pro-attitude in that set. If you like, you can call this the desire-out, desire-in principle.¹⁷ If I can find one transition from belief to pro-attitude and defend the rationality of that transition then we should reject the theory.

While I don't think that the transition from beliefs about what you'll want tomorrow to making preparations today is a fallacy of practical reason, I will focus exclusively on the case of first-person normative belief: beliefs of the form "I ought to ϕ ." Since my knowledge of what sorts of things are morally required and my knowledge of what sorts of things make it the case that things are morally required are both embarrassingly limited, I will focus exclusively on cases in which you think that you ought to ϕ because you think that ϕ -ing is rationally required. If moral requirements turn out to be the same as rational requirements, that will be fine with me. If they turn out to be different, that will be fine as well. The question is whether any beliefs can move the mind to action, so all we need is one.

So we're interested in the rationality of the transition from believing that you ought to ϕ to ϕ -ing. In obvious shorthand, where "B" stands for belief and "O" stands for ought, we're interested in the rationality of (The Move).

(The Move): B[O ϕ]

ϕ -ing

Despite the horizontal line, this is not the claim that ϕ -ing is entailed by believing that you should. This is meant as a description of a transition, a transition that may or may not occur. We're also interested in transitions from the belief that you ought to ϕ to a desire, intention, or attempt to ϕ . Sometimes it will be useful to distinguish the different cases. But they all involve a leap from the theoretical to the practical, so for the most part, we will treat them all the same. And again, to say something good about the transition is not yet to say anything good about the inputs to or outputs

from the transition. Whatever else is wrong with the belief that q , if it comes from the beliefs that p and that if p then q , the problem is not with the transition.

So how should we think about the question of whether (The Move) is a reasonable move for the mind to make? I think it's useful to compare (The Move) to (The Umbrella).

(The Umbrella): B[it's raining]
taking an umbrella

I think that it's fairly clear that the rationality of the transition from the belief that it's raining to taking an umbrella crucially depends on what other mental states you're in. And I think it's safe to assume, at least for the sake of argument, that it crucially depends on what you want. If you want to stay dry and think that an umbrella will help, then this might be a perfectly reasonable transition. But if you don't care about getting wet or don't like umbrellas, the transition might not make sense. If (The Move) depends on what you want in the way the (The Umbrella) seems to depend on what you want, we've made no progress against the Humean theory.

But we can ask about what makes (The Umbrella) make sense in those circumstances in which it does. Moving from the belief that it's raining to taking an umbrella is one important part of what's involved in following a rule, a rule most naturally stated using an imperative that talks about the weather.

(The Rule): If it's raining, take an umbrella.

It's the acceptance of (The Rule) that gives a point to the transition from the belief to the action. And this is still true in cases where we might not say that you're actually following (The Rule). So suppose that you think that it's raining, but you're mistaken. If you take an umbrella as a result, we might not say that you're following (The Rule). But we might say that you're trying to follow (The Rule), and that the point of taking an umbrella, even in this case, is at least partly given by a connection between rain and umbrellas, and not just by a connection between beliefs and umbrellas.

For our purposes, all we need is the following: (The Umbrella), the transition, is at least as rationally acceptable as (The Rule), the imperative. This is not at all meant to suggest that the transition really is the imperative or anything like that. And of course, the rational acceptability of (The Rule) will crucially depend on what you want. Nothing I say should be taken as an argument that all rational agents are required to take their umbrellas. The acceptability of both the transition and the imperative depend on desires and depend on them in much the same way. This is part of the point.

We stated the rule using an imperative, but we could just as easily have used the word “ought.”

(The “Ought”): If it’s raining, you ought to take an umbrella.

I’m not suggesting that the imperative and (The “Ought”) say the same or mean the same or in any sense come to the same thing. But I do think that it’s at least as reasonable to accept (The Rule) as it is to accept (The “Ought”). Maybe there’s more to (The “Ought”) than there is to the imperative. I don’t know. And I’m perfectly happy to assume that the acceptability of this “ought” depends on what you want.

So the transition is at least as acceptable as the imperative which is at least as acceptable as The “Ought.” But the acceptability of all of these things depends on what we want. So how is any of this supposed to help? Well, if there’s an imperative and an “ought” sentence that corresponds to (The Move) in the way that these things correspond to (The Umbrella), then maybe we can evaluate the transition that actually matters to us by evaluating these other things.

So how did we get from (The Umbrella) to (The Rule)? The antecedent of the imperative is the content of the relevant belief, and the consequent is not the claim that you do take an umbrella but an imperative telling you to do so. If we make these changes to (The Move), we get the following.

(Just Do It): If you ought to ϕ then ϕ .

This sounds like good advice to me: if you ought to take an umbrella, then take an umbrella. But since I don’t have a general account of the acceptability of imperatives, let’s move on.

How did we get from (The Rule) to (The “Ought”)? We left the antecedent alone and changed the consequent from an imperative telling you to take an umbrella to a declarative saying that you ought to take an umbrella. If you make these changes to (Just Do It), here’s what you get.

(Gotta Do): If you ought to ϕ then you ought to ϕ .

This is expressed in the vernacular by saying “You gotta do what you gotta do.” And utterances of this apparent tautology have motivational consequences. They’re used to get people to do what they already know that they should.

If (The Move) is at least as acceptable as (Just Do It), which is at least as acceptable as (Gotta Do), then (The Move) is looking pretty good. If (Gotta Do) is, as it appears to be, a tautology, then its acceptability cannot depend on your desires. If p is a tautology, you ought to believe that p whether you want to be reasonable or not, at least when the question comes up. Since

the question of (Gotta Do) has come up, you're rationally required to accept it. So you're rationally required to accept (Just Do It). So (The Move) has the same rational status as the transition from the beliefs that *p* and that if *p* then *q* to the belief that *q*. If the input to the transition is no good, the output may be just as bad. But if (Gotta Do) is a tautology, the transition itself is unimpeachable.

But is (Gotta Do) a tautology? Here there are two ways to worry. You might think that declarative sentences involving the word "ought" do not express propositions, so they can't be either true or false.¹⁸ If you're a non-cognitivist about rationality, you need to give some account of the behavior of "ought" sentences when they figure as antecedents of conditionals.¹⁹ However the details of the account go, you'll need some replacement for truth. The non-propositions expressed by "ought" sentences will not be true, but some of them at least ought to turn out acceptable or required or something like that. And there must be some minimal conditions of adequacy on your account. The point of the account will be to say what's good about transitions like the following:

If you ought to believe that *p*, then *q*

You ought to believe that *p*

So, *q*

If you can't use the notion of truth preservation to say what's good about this, you'd better find something else good to say about it. If you can't, we'll have to reject your theory. And the same goes for (Gotta Do). If it does not turn out to be as acceptable or required as a non-proposition can be, that will constitute sufficient grounds for rejecting the account. So (Gotta Do) is as good as it gets. Either it's a necessarily true proposition, or it's as acceptable or required as a non-proposition can be. Either way, (The Move) is as acceptable or required as a transition can be. So non-cognitivism about rationality, whatever else you'd like to say about such a view, is no threat to (Gotta Do) or (The Move).

Even if we assume that (Gotta Do) expresses a proposition, we might still worry that we have incorrectly identified its logical form if we call it a tautology. Perhaps the second occurrence of "ought" should be given wide scope rather than narrow scope.²⁰ Compare:

(B) If you believe that *p* and that if *p* then *q*, you ought to believe that *q*.

The surface syntax of the sentence suggests that the operator "ought" has narrow scope. It seems to govern the consequent, or a part of the consequent, not the conditional as a whole. We would represent this reading in shorthand as follows:

(NB) $[(Bp \ \& \ B(p \rightarrow q)) \rightarrow OBq]$

This says that you should always accept the consequences of your beliefs no matter how absurd the consequences or unjustified the beliefs. I take it that that's false. Sometimes you just have to change your mind. But there is something true (B) could be used to get across. Just give "ought" wide scope and let it govern the conditional.

(WB) $O[(Bp \ \& \ B(p \rightarrow q)) \rightarrow Bq]$

This says that you ought to make sure the following conditional is true: if you believe a conditional and its antecedent, then you believe the consequent. You can satisfy this obligation in various ways, by believing the consequent or by failing to believe either the conditional or its antecedent.²¹

So if we accept the legitimacy of the distinction between the wide- and narrow-scope "ought," there are two ways of reading (Gotta Do) depending on the scope of the second "ought." Here are the readings in shorthand.

(NGD) $O\phi \rightarrow O\phi$

(WGD) $O[O\phi \rightarrow \text{you } \phi]$

On the first reading, (Gotta Do) is a tautology. The second reading doesn't say that you do everything you're supposed to do. It says that you ought to make sure that you do everything that you're supposed to do. On a fairly minimal reading of the wide-scope "ought," all it does is rule out as impermissible certain combinations. This one rules out as impermissible the following combination: failing to ϕ when you ought to ϕ . I take it as intuitively obvious that that combination is impermissible. If that's what (WGD) says, then it is obviously and perhaps even trivially true.²²

I accept the legitimacy of the distinction between the wide- and narrow-scope "ought." So I think that the difference between the scope of the "ought" usually makes an important difference. It just doesn't make that much difference in this one particular case. On the other hand, some people are suspicious of the distinction between the wide- and narrow-scope "ought." For these people, there's really only one way to read (Gotta Do) and that's as a tautology. So if the distinction collapses, so does the objection to my argument.

So whether we read it with wide or narrow scope; whether we accept or reject the wide/narrow scope distinction; and whether we read it truth evaluably or otherwise, (Gotta Do) is as good as it gets. So (The Move) is as rationally required as any transition can be. So suppose that you believe that you ought to ϕ . The desire to do what you should is not rationally required to produce a desire, intention, or attempt to ϕ . And the desire is not, at least

as far as we can tell from the armchair, causally required to produce the relevant pro-attitudes or action. If there's any sense at all in which the desire really is necessary, that sense remains to be explained. So either the Humean Theory of Motivation is false, or we don't know what it means.

The Deep Argument

I don't think there's anything wrong with the Quick and Easy Argument. I think you can derive the negation of the Humean theory from a tautology. Though the popularity of the Humean theory is waning, especially among the ethicists,²³ many people have found it a very natural view. So you might think that a deep question like whether reason alone can ever move the mind deserves a deep argument to answer it. And you might think that while the Quick and Easy Argument should convince us that (The Move) is a reasonable move for the mind to make, it doesn't really explain why it's a reasonable move. So I have another argument in favor of the rationality of (The Move). Perhaps "deep" is not the right word for this argument. Perhaps there's some other word that means "there seems to be something to this argument, but I can't quite figure out what." In any case, first I'll state the premises, then derive the rationality of (The Move) from them, and finally do what I can to make the premises seem plausible.

Consider this claim about what your actions commit you to.

(AC) If you intentionally ϕ , you're (rationally) committed to its being (rationally) permissible for you to ϕ in those circumstances.

This is an expression of the idea that you're rationally required to be reasonable, not just in thought but in action as well. If it's rationally permissible for you to do something, then it's a reasonable thing for you to do. And if you do something on purpose, the question of whether it was reasonable can always come up. If the verdict goes against you, you can always say, "But I never said I was reasonable." This can be a good joke. But that's the most it can be.

Now consider this claim about your doxastic commitments.

(DC) If you believe that p , you're (rationally) committed to p 's being true.

To say that you're committed to something's being true or reasonable is not to say that you ought to believe that it's true or reasonable. If you believe the wrong things, you could end up committed to a contradiction. But that doesn't mean that you ought to believe it. This is one of those cases where you ought to change your mind. And if you're committed to something that turns out to be false, that's a problem for you even if you believe it. So you're committed as it were to it, not just to believing it.

Suppose you intentionally ϕ when you think that you shouldn't ϕ or that ϕ -ing is rationally impermissible or unreasonable. The (AC/DC) combination entails that you are committed to its being reasonable to ϕ (by AC) and to its being unreasonable to ϕ (by DC). Your action conflicts with your beliefs in a straightforward way. Together they commit you to a contradiction. Being committed to a contradiction reflects on your rationality (and not in a good way).

Now, (AC/DC) doesn't tell you not to ϕ in the relevant situation, and it doesn't tell you to stop believing that ϕ -ing is unreasonable. All it does is rule out certain combinations. In this respect it resembles (WB).

(WB) $O[(Bp \ \& \ B(p \rightarrow q)) \rightarrow Bq]$

This says that you shouldn't believe a conditional and its antecedent without also believing the consequent. Sometimes you should revise, and sometimes you should conclude. But (WB), on its own, cannot tell you which.

So which combinations does (AC/DC) rule out? Well, it rules out believing that you shouldn't ϕ while intentionally doing it anyway. And presumably, it rules out believing that you shouldn't refrain from ϕ -ing while intentionally refraining anyway. But if you shouldn't refrain from ϕ -ing then you ought to ϕ . So (AC/DC) rules out believing that you ought to ϕ while intentionally failing to ϕ . Now suppose that you think you ought to ϕ because you think that ϕ -ing is the most reasonable thing to do in the circumstances; you try to ϕ ; but through no fault of your own, you fail. We might think that you need to hone your skills, but we won't think that you're being unreasonable. You fail to do what you think you should. But you don't intentionally fail to do what you think you should. So if we restrict the variable " ϕ " to those things you're in a position to do intentionally, or those things you're in a position to do for a reason, it looks like (AC/DC) entails the following wide-scope "ought."

(WSO) $O[BO\phi \rightarrow \text{you } \phi]$

This says that you ought to make the following conditional true: if you think you ought to do something then you do it. You can fulfill this obligation either by lacking the belief or by performing the action.

But what does any of this have to do with (The Move)? To say that (The Move) is a rationally required transition is not to say that you ought to ϕ every time you think you should. If the inputs to the transition are no good, the output might be just as bad. To say that the transition is required is just to rule out certain combinations: believing you ought to ϕ without intending or trying to ϕ . And if ϕ -ing is something you're in a position to do on purpose, it rules out believing you should without doing it. So (AC/DC) rules out exactly the same combinations as (WSO) which rules out the same

combinations as (The Move). If those combinations are in fact rationally prohibited, then we have both an argument that and an explanation of why (The Move) is rationally required.

So we can get to the rationality of (The Move) from some claims about commitment. Is there any reason to think those claims are true? I take it that (DC) is fairly clear, possibly even more clear than the idea of one thing committing you to another. Whatever “committed to” means, (DC) better turn out to be true. So what reason do we have to believe (AC)? I kick you in the shins. Why does that commit me to its being reasonable for me to do so? If I intentionally kick you in the shins, then I kick you in the shins on the basis of a reason, a merely motivating reason to be sure, but a reason none the less, and not just any old reason. If I intentionally kick you in the shins, I act on a reason for kicking you in the shins.

Now I might have (motivating) reasons for kicking you in the shins running around in my head without taking them at all seriously. I might even try to get rid of them. But if I act on those reasons, then I’m taking them seriously. I’m endorsing them, and I’m responsible for them. So what is it to endorse or be responsible for a reason to ϕ ? Well, what’s a reason to ϕ ? A reason to ϕ is something that’s supposed to make ϕ -ing reasonable, and good reasons are things that do what they’re supposed to do. They’re things that make things reasonable. If you endorse something whose job it is to make ϕ -ing reasonable, it looks like you’re committed to its doing its job. You’re committed to the rationality of ϕ -ing. So if you do something intentionally, you’re committed to its being a reasonable thing to do.

The Consequences

The rejection of the Humean Theory of Motivation is certainly not new. Nor is the acceptance of (The Move). If there’s anything new so far, it’s in the arguments for (The Move). But while a number of people have accepted it, I don’t think that the consequences of accepting (The Move) have been fully appreciated. There are two separable issues about the role of the mental in the determination of practical reason. We can ask about the beliefs, or we can ask about the desires, and so far, we’ve been focusing primarily on the desires. But the legitimacy of (The Move) also has consequences for the role of belief.

As we’ve seen, accepting (The Move) is as reasonable as accepting the following wide-scope “ought,” at least when “ ϕ ” is restricted to things you’re in a position to do for a reason.

(WSO) $O[BO\phi \rightarrow \text{you } \phi]$

From this, plus the fact that you believe that you ought to ϕ , it doesn’t follow that you ought to ϕ . This inference is called “detachment,” and it’s generally

agreed that it's invalid.²⁴ (WSO) only rules out certain combinations: it says that you shouldn't believe that you should without doing it if you can. From the mere fact that you think you should, all that follows is that you should either revise that belief or perform the action.

But now consider what we might call "strengthened detachment."

(1) $O[BO\phi \rightarrow \text{you } \phi]$

(2) $OBO\phi$

(3) $O\phi$

Here we add not just that you do believe but that you should believe. The first premise says that you ought to either give up the belief or perform the action. The second premise says that the first option is not permissible. If you ought to believe that p then failing to believe is not permissible. So the only permissible option left is to perform the action. If there's only one permissible option, that's what you ought to do. So if the fans of the wide-scope "ought" are right about what (WSO) means, i.e., that you ought to either revise the belief or perform the action, it looks as though strengthened detachment is valid.

If (1) and (2) entail (3), then (1) by itself entails the conditional from (2) to (3).

(23) $OBO\phi \rightarrow O\phi$

If you ought to believe that you ought to ϕ , then you ought to ϕ . What are some circumstances in which you ought to believe something? Suppose that you're in a situation in which you have to do one of the following three things: either believe that p , withhold judgment on p , or deny that p . If withholding judgment is simply failing to either believe or deny, you are always in such a situation with respect to every proposition. If the proposition that p has never occurred to you and there's no reason why it should, then withholding judgment in this very weak sense is the way to go. Now suppose further that with respect to some particular proposition, it's more reasonable for you to believe that p than it is to either withhold or deny. What should you do? You should be reasonable. You should believe that p .

When we say that you are justified in believing that p , we either mean that your belief that p is based on sufficient evidence, or that there is sufficient evidence in favor of believing that p available to you whether or not you believe or believe for the right reasons. But how much evidence is sufficient? Presumably, if it's more reasonable for you to believe that p than it is for you to either withhold or deny, then that would be sufficient. So if we think of being justified in terms of having this degree of justification, and we let "J" represent this idea, we can rewrite (23) as follows.

(JO) $JO\phi \rightarrow O\phi$

If you're justified in believing that you ought to ϕ , then you ought to ϕ . I think that (JO) is true and important and that it causes trouble for any attempt to understand practical reasons in terms of the facts instead of in terms of the beliefs.

(JO) has an analogue in the theoretical realm.

(JJ) $JJp \rightarrow Jp$

If you're justified in believing that you're justified in believing that p , then you're justified in believing that p .²⁵ For most contingent propositions, justification does not entail truth. But according to (JO) and (JJ), propositions about what's currently most reasonable for you to think or do provide exceptions to this general rule. And in both cases the basic idea is the same: you can't make a non-culpable mistake about what's reasonable for you to do in your current circumstances. Of course, you can make a mistake about what you ought to do. If believing you should made it true that you should, then "should" wouldn't be worth the ink it takes to write. But that's not what (JO) or (JJ) say. They say, in effect, that you're responsible for the facts about what's reasonable for you to do or think.

So suppose that you think that you ought to ϕ in circumstances in which you really shouldn't. According to (JO), this is a culpable error on your part, and your belief is not justified. You should have known better. This is a fairly serious epistemic burden on the rational agent. But it's an expression of the idea that it is the duty of the rational agent to be rational. We don't merely expect you to try to be reasonable or to do what you think might be reasonable in the circumstances. We expect you to succeed in actually being reasonable. Of course, occasional lapses are explicable and understandable. But I wouldn't have thought that failures of rationality are rationally permissible. So if you have to be reasonable, then you have to know what would constitute being reasonable in the circumstances in which you find yourself. So you're responsible not only for acting in accord with the dictates of reason. You're responsible for knowing what those dictates are. So mistakes about what's reasonable are automatically the agent's fault in a way that not all false beliefs are automatically the agent's fault.

So (JO) places fairly serious constraints on the rational agent: (JO) expects you to be reasonable. But the constraints are not only on the agent. (JO) also places fairly serious constraints on what sort of facts can make it the case that you ought to ϕ . If the facts that made it the case that you ought to ϕ were completely inaccessible to you, then you'd have no way of knowing that you ought to ϕ . You might even be justified in believing that you shouldn't. So if you fail to do what you ought to do in this case, we couldn't hold that failure against you. But if we can't hold a failure against a rational agent,

that failure cannot be a failure of rationality. So at least when we're talking about what you ought to do in the sense of what you're rationally required to do, the facts that make it the case that you should, or the fact that you should, must be accessible to the agent. And we can use (JO) as a test to determine just how accessible they have to be.

So return to Smith's story about the Picasso. The fact that it's a Picasso is supposed to be a normative reason for him to buy it, given, of course, a desire to buy a Picasso. He doesn't know it's a Picasso, and there's nothing in the story to suggest that he should have known. So he's at least justified in withholding judgment on the proposition that it's a Picasso. So he's justified in believing that he shouldn't buy the painting, at least until he finds out who painted it. So if he does what he supposedly has a normative reason to do, namely buy the painting, he'll be acting against his better judgment: he'll be buying while justified in believing that he shouldn't. So if you accept (The Move), you must reject not only the Humean Theory of Motivation. You must also reject Smith's conception of normative reasons. Either the fact that it's a Picasso is not a normative reason, or normative reasons don't make it the case that you ought to do one thing rather than another.

I take it that Smith's conception of normative reasons is fairly standard. Bernard Williams tells a story in which what you think is a gin and tonic in fact contains petrol.²⁶ He says you have no reason to drink. The reasons, in effect, go with the facts, not the evidence. Derek Parfit imagines a case in which he falsely believes that his hotel is on fire.²⁷ While he concedes that this might make it rational for him to jump, he doesn't think it provides a normative reason to jump. But if he's justified in believing that his hotel is on fire, then he's justified in believing that he ought to jump. So at least according to (JO), that's what he ought to do. And unlike Smith and Williams, Parfit does seem to accept (The Move).

Jonathan Dancy seems to have the same basic conception of normative reasons, despite his requirement that the objective facts that constitute normative reasons must be capable of passing through "the agent-relative epistemic filter."²⁸ Basically, this means that the relevant facts must be accessible to the agent in some sense. Consider the following story and Dancy's reaction to it.

Suppose, for instance, I make a not very sensible choice about what arrangements to make about my pension. And suppose that I can later explain the choice I made by pointing out that there were some crucial facts that I happened quite reasonably to have got wrong, and in this way, as we might put it, exculpate myself. There is a sense of 'justify' in which I can be said to have justified doing what I did. But this does not show that the balance of reasons was in favour of the action. It wasn't. Indeed, the features in light of which I made my choice turned out all to be a mistake, and so cannot count even as defeated reasons.²⁹

Dancy “quite reasonably” gets the facts about his pension wrong. So he has justified false beliefs about his pension. Since he can justify his choice on the basis of the beliefs, there must not be anything wrong with the transition from those beliefs to the belief that he ought to make the arrangements. So he’s justified in believing that he ought to make the arrangements. This justified belief cannot make the arrangements the most profitable, or most beneficial, and it cannot make them produce the greatest happiness for the greatest number. But at least according to (JO), it does make the choice the most reasonable. So if the balance of some kind of normative reasons favors some other option, then those reasons favor being unreasonable. I can’t imagine what kind of reasons those could be.

I don’t know in what sense the facts about the pension were supposed to be accessible to Dancy. But they weren’t accessible in the sense that he was justified in believing them. So the degree of accessibility required by (JO) is much more serious than that required by Dancy. And if your reasons and what you ought to do must be grounded in the objective facts rather than in the agent’s take on the facts, then the story constitutes a case in which the agent is justified in believing that he ought to ϕ even though he really shouldn’t. So you can’t have both (JO) and this conception of normative reasons. That means that you can’t have this conception of reasons and accept the rationality of (The Move).

Does this mean that your practical obligations are completely determined by your non-factive mental states like beliefs and desires because those are the only things we have access to? Absolutely not. You have access to and justified beliefs about all sorts of facts about the external world, from what you had for breakfast to whether or not it will hurt if you kick me in the shin. And if you ought to know that p , because the fact that p is staring you in the face, but you’re too careless to notice that p , there may be nothing in your inner life or your non-factive mental states that makes it the case that you have available all the evidence you need to reasonably believe that p . But the claim that you ought to know that p , like the claim that you’re in a position to know that p , is a claim about your perspective or your take on things, despite the fact that both these claims entail that p . These are the sorts of things that make things reasonable. And they’re the sorts of things that make it the case that you ought to do what you ought to do.³⁰

Notes

¹ Smith (1994): 94.

² Ibid p. 92.

³ Ibid p. 94.

⁴ Ibid p. 97.

⁵ Ibid p. 95.

⁶ Ibid p. 95.

⁷ Davidson (1980).

⁸ Though this picture doesn't fit with some of the things Smith says, it fits quite nicely with other things he says. In both Pettit and Smith (1990) and Smith forthcoming for example, we're told that motivating reasons don't just explain actions, they rationalize action, where I think that means that they determine the rationality of action in the way I describe. And in the latter paper, the analogy with the theoretical case is explicit.

⁹ Nagel (1970): 39–40.

¹⁰ Smith (1994): 100.

¹¹ *Ibid.* p. 116.

¹² *Ibid.* p. 115.

¹³ Anscombe (1957).

¹⁴ *Ibid.* p. 56.

¹⁵ This is Anscombe's example, p 56.

¹⁶ For some serious attempts, see Velleman (1992) and Humberstone (1992).

¹⁷ Wallace (1990).

¹⁸ For versions of non-cognitivism about moral discourse, see Ayer (1946), and Blackburn (1984). For non-cognitivism about notions like justification and being reasonable, see Gibbard (1990) and Field (2000).

¹⁹ This is a familiar problem for non-cognitivism. See Geach (1965).

²⁰ I'd like to thank Sruthi and Cullen for this suggestion. For more on the difference between and the significance of the difference between wide- and narrow-scope "oughts," see Broome (1999).

²¹ I'm assuming that if the beliefs are yours then so is the relevant obligation.

²² I take it that the problem with the fairly minimal reading of the wide-scope "ought" is that on that reading, the following are both true (at least, the first one's true if you ought to brush your teeth).

O[Grass is red or grass is not red → You brush your teeth]

(BAD)O[You do something wrong → You make sure someone else gets blamed]

The first of these comes from Broome (1999) and the second from an anonymous referee, though Broome thinks his is true and the referee thinks the second is false. There's certainly something peculiar about both of these. The tautology does not require you to brush your teeth. There's no normative connection at all between the tautology and your brushing your teeth. In the second case, doing something wrong doesn't require you to blame someone else. If anything, it requires you to make sure that no one else gets blamed. So there may be a normative connection between antecedent and consequent, but certainly not the one suggested by (BAD).

Nevertheless, on the minimal reading of the wide-scope "ought," both of these come out true. The following combination is impermissible: the tautology is true and you fail to brush your teeth. The combination is impermissible, not because of any connection between the two, but simply because you ought to brush your teeth. Similarly, the following combination is also impermissible: you do something wrong and don't get anyone else in trouble for it. The combination is impermissible not because you should have gotten someone else in trouble, but simply because you shouldn't have done anything wrong in the first place.

Contrast these with the usual examples of the wide-scope "ought." You can see that *p* obviously entails *q*. So believing *p* requires believing *q*. You're not supposed to imagine the case in which *p* is a contradiction so it entails everything and there is always something wrong with any combination that includes believing *p*, no matter what else it includes. You're supposed to imagine the case in which *p* has something to do with *q*, so there's a non-arbitrary normative connection between believing *p* and believing *q*. Even without an analysis or theory about which normative connections count as non-arbitrary, I take it as intuitively obvious that being required to *f* requires you to *f*. If the connection between being required to *f* and *f*-ing doesn't count as a non-arbitrary normative connection, then probably nothing does, and you should seriously doubt the significance of the more substantive reading of the wide-scope "ought." So whether we read the wide-scope "ought" minimally or more substantively, (WGD) is necessarily true.

I'd like to thank an anonymous referee for *Nous* not just for this example but for other help on this section as well.

²³ See, for example, Nagel (1970), McDowell (1978), Platts (1979), and Parfit (1997).

²⁴ At least among those who accept the legitimacy of the distinction between the wide- and narrow-scope "ought." See Broome (1999).

²⁵ I defend this in Gibbons (2006).

²⁶ Williams (1980): 102.

²⁷ Parfit (1997): 99.

²⁸ Dancy (2000): 56–57.

²⁹ *Ibid.* p. 7. Also see pp 62–63 for a structurally similar story about a boy with a justified false belief about a snake in his sleeping bag.

³⁰ I would like to thank David Henderson, Joe Mendola, David Sobel, Mark van Roojen, and an anonymous referee for *Nous* for helpful comments on earlier versions of this paper.

References

- Anscombe, G.E.M. (1957), *Intention* (Oxford: Blackwell).
- Ayer, A.J. (1946), *Language, truth and Logic* 2nd ed. (Harmondsworth: Penguin).
- Blackburn, Simon (1984), *Spreading the Word* (Oxford: Oxford University Press).
- Broome, John (1999), "Normative Requirements," *Ratio* 12: 398–419.
- Dancy, Jonathan (2000), *Practical Reality* (Oxford: Oxford University Press).
- Davidson, Donald (1980), "Actions, Reasons, and Causes," in *Essays on Actions and Events* (Oxford: Clarendon).
- Field, Hartry (2000), "Apriority as an Evaluative Notion," in *New Essays on the A Priori* eds. Paul Boghossian and Christopher Peacocke (Oxford: Oxford University Press).
- Geach, P.T. (1965), "Assertion," *Philosophical Review* 74: 449–465.
- Gibbard, Allan (1990), *Wise Choices, Apt Feelings* (Oxford: Clarendon).
- Gibbons, John (2006), "Access Externalism," *Mind* 115: 19–39.
- Humberstone, Lloyd (1992), "Direction of Fit," *Mind* 101: 59–83.
- McDowell, John (1978), "Are Moral Requirements Hypothetical Imperatives?" *Proceedings of the Aristotelian Society* supp. vol. 52: 13–29.
- Nagel, Thomas (1970), *The Possibility of Altruism* (Princeton: Princeton University Press).
- Parfit, Derek (1997), "Reasons and Motivation," *Proceedings of the Aristotelian Society* supp vol 71: 99–130.
- Pettit, Philip and Michael Smith, "Backgrounding Desire," *Philosophical Review* 99: 565–592.
- Platts, Mark (1979), *Ways of Meaning* (London: Routledge and Kegan Paul).
- Smith, Michael (1994), *The Moral Problem* (Oxford: Blackwell).
- forthcoming, "Is There a Nexus Between Reasons and Rationality?" in *Moral Psychology* ed, Sergio Tenenbaum (Amsterdam: Rodopi).
- Velleman, J. David (1992), "The Guise of the Good," *Nous* 26: 3–26.
- Wallace, Jay (1990), "How to Argue about Practical Reason," *Mind* 99: 267–297.
- Williams, Bernard (1980), "Internal and External Reasons," in *Moral Luck* (Cambridge: Cambridge University Press).

Copyright of *Nous* is the property of Blackwell Publishing Limited and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.