

Things That Make Things Reasonable

JOHN GIBBONS

University of Nebraska

There are many things to ask about reasons. We can ask in some particular circumstances what reasons you have for going to the store or for believing that you're out of milk. But we can also ask, much more generally and abstractly, whether or not the conditions under which someone has a reason to ϕ are determined that person's perspective or point of view. I take it that if you know, or are in a position to know, or are justified in believing, or believe that you're out of milk, these are all facts about your perspective. The fact that you're out of milk, on the other hand, is not a fact about your perspective. Given the relatively light work I'm going to ask the notion of a perspective to do, I take it that the idea is clear enough.

It would be very odd if it turned out that practical and theoretical reasons differed from each other at this level of abstraction. Reasons to believe are one way, but reasons to act are another. I'm not sure it's impossible for practical and theoretical reasons to differ in this way. It's just that similarities between practical and theoretical reasons have a built-in explanation: they're both reasons. Purported differences between practical and theoretical reasons stand in need of explanation in a way that similarities don't. And the best explanation for such a difference would ultimately explain it in terms of the fact that these reasons are practical and those reasons are theoretical. What's impossible for me is to see how the difference between the theoretical and practical, whatever that difference may be, could result in or explain the difference between perspective-dependence and perspective-independence.

We can ask questions at various levels of abstraction about the conditions under which it would be reasonable or rational for someone to ϕ . Again, we can ask whether or not these conditions are determined by the person's perspective. And again, it would be odd, though

perhaps not impossible, for our answers to diverge in the practical and theoretical cases. This is not just a penchant for symmetry. There must be something to the idea that if it's reasonable to believe that p then it's reasonable to act as though p . But if our answers to these questions do diverge, and if there is no substantive explanation of why, then we should seriously doubt that there is a single, coherent notion of a reason or of rationality that covers both the practical and theoretical cases.

In addition to thinking about the relation between our views on practical and theoretical reasons, or practical and theoretical rationality, we also need to think about the relation between our views on reasons and our views on what's reasonable. It certainly seems, at least from these lofty heights, that the notion of a reason and the notion of what's reasonable are in the same normative ballpark. So similarities with respect to very general features like perspective-dependence or -independence should not come as a surprise. But differences in these respects need an explanation, ultimately in terms of the fact that this is about a reason but that's about what's reasonable.

We can't always get what we want. But I think deep down, what we'd like best is either an account of these things that made them all perspective-dependent, or an account that made them all perspective-independent, or at the very least, an extremely good explanation of why things fall apart. Let's see what we can get.

To descend from our lofty heights, let me tell you some stories and ask you some questions. Last evening, you checked and made sure that you had enough milk for breakfast. Late last night, a professional dairy thief broke into your house, stole your milk, and left no trace: no broken windows, screaming burglar alarms, or anything like that. This morning, you're standing in the kitchen, but you have not yet opened the refrigerator door. You have no milk but believe that you do. Do you have any reason to believe that you're out of milk? And would believing that you're out of milk be reasonable in these circumstances?

I expect fairly general agreement on the answers to these questions. No doubt, there are other circumstances in which such a belief would be perfectly reasonable, for example, if you opened the refrigerator door. But in these circumstances, it seems that there's nothing to make the belief reasonable. And it also seems that you have no reason; there is no reason; there's no such thing as your reason; and the same goes for any other way of saying the same thing. In this one particular case, it seems that reasons and rationality go together and that both are perspective-dependent. Whether a belief is reasonable or you have reason to believe doesn't just depend on the facts. It depends on the evidence available to you.

You're standing in the kitchen, but you have not yet opened the refrigerator door. You think you have milk; you don't have milk; and you want some. Do you have any reason to go to the store to buy milk? And would it be reasonable for you to go to the store for milk in these circumstances? I think that our answers to these two questions should be the same and that they should be the same as our answers to the analogous questions about theoretical reasons and rationality. And this is exactly what my intuitions tell me. Going to the store for milk would not be the least bit reasonable in these circumstances. No doubt, there are other circumstances in which that would be perfectly reasonable, for example, if you opened the refrigerator door. And no doubt, if you went to the store and bought milk by mistake, that would be a very good thing. But buying milk by mistake is not a reasonable thing to do. In these circumstances, there's nothing to make the action reasonable. So you have no reason; there is no reason; and so on. These are my intuitions, and while I know I'm not alone in this, I do not expect fairly general agreement on these answers.

The other story is similar, but sometimes people have slightly different reactions depending on how the facts come apart from the perspective. So we should look at this one as well. This time, you checked for milk last night, discovered that there wasn't any left, and planned to go to the store in the morning. Late last night, a professional dairy anti-thief broke into your house, left some milk in your refrigerator, but otherwise left no trace. This morning, you're standing in the kitchen, but you have not yet opened the refrigerator door. You have milk. But do you have any reason to believe that? And would such a belief be reasonable if it were formed in these circumstances? Again, there are other circumstances in which the belief would be reasonable. But in these circumstances, there's nothing to make it reasonable. So you have no reason to believe.

You're standing in the kitchen, but you have not yet opened the refrigerator door. You believe you're out of milk; you're mistaken about that; and you don't need anything else at the store. Do you have a reason to go to the store to buy milk? Would going to the store for milk be reasonable in these circumstances? I take it as fairly obvious that going to the store for milk in these circumstances would be a reasonable thing to do and expect fairly general agreement on this. But as we'll see, it turns out to be a somewhat complicated question whether this expectation is met. I also think that in this case you have a reason to go to the store, by which I mean that you have a *good* reason to go to the store. You have the kind of reason to go to the store that is intimately related to some positive normative status. You have the kind of reason to go to the store that makes it reasonable to go to the store.

If being reasonable is not a positive normative status, I don't know what is.

Surprisingly enough, to me at least, there are people who deny that you have a reason, or at any rate a good reason, to go to the store in this case when you're justified in believing that you're out of milk. Sometimes the idea is that there are two kinds of reasons, and while you may have a reason of one kind, you don't have a reason of the other. Quite often, these people hold a similar view in the dairy thief case. When you mistakenly think that you already have milk, you have reason to go to the store. The basic idea, developed in different ways by different people, is that practical reasons (or practical reasons of a certain kind) are perspective-independent. Just for the sake of a name, let's call these people "the ethicists."¹ There is, to put it lightly, a certain amount of disagreement among the ethicists about exactly what role your desires play in the determination of your practical reasons. But the people I'm talking about all agree that the relevant practical reasons are independent of what you believe, or are justified in believing, or what you're in a position to know. This is the only kind of perspective-independence that I'm talking about in this paper.

The idea that practical reasons are perspective-independent can be combined with the idea that practical rationality is perspective-dependent. Such a view requires an explanation of why reasons go one way while being reasonable goes the other way, but it's certainly a view to consider. But the idea that practical reasons are perspective-independent can instead be combined with the idea that practical rationality is also perspective-independent. This view is committed to the wildly implausible claim that if you go to the store for milk when you think you don't need any, this is not just lucky or a good thing. This is a reasonable thing to do. Though the view is obviously false, when we look at some things that ethicists actually say, we'll see why we need to keep the view on the table.

So there are three views about reasons and rationality to consider. For purely abbreviatory purposes, so we don't have to state each view every time we discuss it, I'll adopt the following convenient labels.

The Good: Both reasons and rationality are perspective-dependent.

The Bad: Reasons are perspective-independent but rationality is perspective-dependent.

¹ See, for example, Collins (1997), Dancy (2000), Hyman (1999), Parfit (1997), Raz (2005), Smith (1994), Thomson ((2003), Wiland (2003), and Williams (1980) and (1989).

The Ugly: Both reasons and rationality are perspective-independent.

I take it that The Good is the correct account of the theoretical case,² so our concern will be with the practical realm. I also take it that no one (almost no one?) accepts The Ugly on purpose. Its role in the debate is primarily cautionary: if you're not careful, you could end up committed to it. I think that The Bad, with its fractured conception of reasons and rationality is inherently unstable. To the extent that you can't let go of the idea that reasons are perspective-independent, this instability will drive you to The Ugly. To the extent that you're moved by the idea that what's reasonable for an agent to do is determined by the agent's point of view, you'll be driven into the arms of The Good. This is exactly where I hope to drive you.

In addition to asking, in a very general way, about the conditions under which you have a reason to ϕ , we can also ask about what kind of thing, metaphysically speaking, a practical reason is. If you accept The Good, it's quite natural, though perhaps not mandatory, to accept what's known as psychologism. This is the view that practical reasons are psychological states or facts about your psychological states. Given that I've framed the issue in terms of perspective-dependence and – independence, you might prefer perspectivism to avoid having to argue about what does and what does not count as a psychological state. Perspectivism is the view that practical reasons are or are facts about elements of your perspective. I take it as obvious that the fact that you know that p is a fact about your perspective, but not everyone thinks that knowledge is a mental state. But since “psychologism” is already in the literature, and since I think that knowledge is a psychological state,³ I'm happy to adopt the title “psychologist.”

If you end up accepting The Ugly, it's natural to accept objectivism about practical reasons: reasons are perspective-independent facts. The fact that you're out of milk not only gives you a reason to go to the store whether you know it or not. It makes it reasonable for you to go to the store whether you know it or not. People who accept The Bad customarily adopt dualism about practical reasons. Typically the idea is that some reasons, motivating reasons, are determined by, and

² I myself doubt this, but there may be a few basic a priori propositions you ought to believe no matter what evidence you have. If there are such things, they'll be perspective-independent in our sense, and our question becomes whether, in the overwhelming majority of cases, practical reasons and rationality are perspective-dependent. If the practical realm is analogous to the theoretical, there will be at most a few things you ought to do no matter what, and I take it that going to the store for milk will not be among them.

³ For arguments for this view, see Williamson (1994), McDowell (1995), and Gibbons (2001).

probably are your psychological states. But the other reasons, the so-called normative reasons, are determined by and probably are the objective facts. So the standard Bad line on the dairy thief case would be to say that you have no motivating reason to go to the store, but you do have a normative reason.

The Ugly

I suggested that when you look at some of the things ethicists say, it's not completely clear that they reject The Ugly. So let's look at some things some ethicists say. Bernard Williams considers a case in which there's some petrol in a glass, but the agent thinks it's gin and wants a gin and tonic.⁴ Does the agent have a reason to drink what's in the glass? Williams says no. After presenting a number of arguments against his own view (e.g., if the agent drinks he'll drink for a reason), Williams rejects these arguments on the basis of the following consideration.

[Saying that the agent has a reason to drink] looks in the wrong direction, by implying in effect that the [resulting conception of reasons] is only concerned with explanation, and not at all with the agent's rationality, and this may help to motivate a search for other sorts of reasons which are concerned with his rationality. But [Williams' conception of reasons] is concerned with rationality.⁵

So Williams' conception of reasons is concerned with the agent's rationality, and on that conception, the agent has no reason to drink what he thinks is a gin and tonic. This probably doesn't entail but it does at least suggest that doing what you have no reason to do would be unreasonable. But if, on the other hand, drinking what you take to be a gin and tonic when you want one is a reasonable thing to do, and if Williams' conception of a reason says that you have no reason to drink, then I, for one, am completely motivated to search for another conception of reasons, one that is, let's just say, *more* concerned with the agent's rationality. I would go looking for the kind of reason for A-ing whose presence is somehow connected to its being reasonable to A, rather than the kind of reason for A-ing whose presence or absence is independent of whether or not A-ing is reasonable.

Williams never comes right out and says that every case of acting on a false belief is irrational. But things he says suggest it. This is not an isolated incident, and it has nothing to do with whether you

⁴ Williams (1980): 102.

⁵ Ibid pp. 102-103

are what the ethicists call an internalist or externalist about reasons.⁶ Michael Smith considers a case in which he wants to buy a Picasso; the painting in front of him is a Picasso; but he doesn't know this.⁷ In Williams' case, you have the belief that "speaks in favor of" drinking, but you don't have the fact that speaks in favor of drinking. This is like our anti-thief case. In Smith's case, you have the fact that speaks in favor of buying, but not the belief. This is like our dairy thief case. After saying that in his case he has a reason to buy the painting, he goes on to say that there is a "requirement of rationality" that he buy the painting, and slightly later that he is "rationally required" to buy it.⁸

I don't have anything like a complete list of what rationality requires. But I would have thought that rationality requires you to do reasonable things and not do unreasonable things. I would have thought that if ϕ -ing is rationally required then it's rationally permissible, and if it's rationally permissible, then it's a reasonable thing to do. Smith does not come out and say that buying the relevant painting (think about how much a Picasso costs) without knowing it's a Picasso or even being all that interested in it for any other reason, is a reasonable thing to do. But things he says suggest it.

Finally, Derek Parfit considers a case in which he falsely believes that his hotel is on fire.⁹ I think we can safely add to the story that he has no other reason to jump out the window, say, a sudden urge to go swimming in the canal below. This makes it like the anti-thief case. Since Parfit thinks that "[w]hile reasons are provided by the facts, the rationality of our desires and acts depends instead on what we believe,"¹⁰ it's natural to suppose that he thinks you have no reason to jump, even though jumping would be rational. The problem with this interpretation is that immediately after saying that he may have no reason to jump even though jumping may be rational, he goes on to say "to be rational is to respond to reasons."¹¹

The following cannot all be true on the most natural interpretation.

⁶ Though there are many things in the debate about practical reason called "internalism," most of these are about the relation between reasons and motivation and have nothing to do with the question of whether or not reasons are perspective-dependent in the sense at issue.

⁷ Smith (1994): 94.

⁸ Ibid p. 97

⁹ Parfit (1997): 99.

¹⁰ Ibid p. 99

¹¹ Ibid p. 99

- (1) He has no reason to jump.
- (2) It's rational for him to jump.
- (3) To be rational is to respond to reasons.

If (2) it's rational for him to jump and (3) being rational is responding to reasons then (not-1) there must be some reasons for him to respond to. And if (3) being rational is responding to reasons and (1) he has no reason to jump, then jumping out of a building you believe to be on fire is irrational. Since the only relevant feature of the belief is that it's false (he doesn't assume that the belief is irrational or unjustified), we're back to the idea that any case of acting on a false belief is irrational. The problem is not that Parfit believes this. He clearly thinks that it's false. The problem is that he seems to be committed to it.¹²

The problem for each of our theorists is exactly the same in each case. Once you drive a wedge between reasons and rationality, none of the ordinary, commonsense things that you want to say about the relation between reasons and rationality will be compatible with your theory. Of course, being reasonable is responding to reasons. Of course, to the extent that you're reasonable you'll do what you have most reason to do. Of course, you're rationally required to be reasonable. Of course, if you have all the reason in the world to ϕ and no reason not to, you're rationally required to ϕ . Of course, if the reasons in favor of ϕ -ing outweigh the reasons against, ϕ -ing is more reasonable than not ϕ -ing. Of course, if you have a reason to ϕ , it's prima facie reasonable for you to ϕ . And of course, a conception of reasons must be concerned with the rationality of the agent. It's not that these platitudes are false. They're just inconsistent with a certain view of reasons.

If you try to combine the platitudes with the idea that reasons are perspective-independent, then you end up with The Ugly, and you must reject (2). If you combine the platitudes with the idea that rationality is perspective-dependent, you end up with The Good, and you must reject (1). But if you think that reasons are perspective-independent while rationality is perspective dependent, then you get The Bad, and you

¹² Parfit (ms.) clearly rejects (3). There the idea is that rationality is responding to reasons or apparent reasons (see the beginning of Chapter 2). So being reasonable is acting on your beliefs about normative reasons. Scanlon (1998) is also inclined toward this view. This seems to have the minor problem that people without thoughts about normative reasons (which may be more of us than you think) are incapable of rationality. But the main difficulty is that since you can have both reasonable ignorance and justified false belief about normative reasons, this just is the wedge between reasons and rationality that's definitive of The Bad. And one of the main purposes of this paper is to explain what's wrong with that.

must reject (3) and the rest of the platitudes along with it. Though some things they say may suggest otherwise, I take it that Williams, Smith, and Parfit all agree on The Bad. This view deserves serious consideration.

The Bad

Whenever an interesting notion has two sides to it and we can't quite see how the two sides fit together, there will always be a temptation to say that there are really two notions. The notion of a reason for action clearly has two sides to it. If you go to the store for milk in the anti-thief case, you will go there for a reason. So you must have a reason. So there must be a reason. But furthermore, we can explain your going to the store in terms of your reasons, and such an explanation will be quite different in interesting ways from an explanation of an explosion in terms of a faulty valve. This is not to suggest that reason-giving explanations are not causal. It's only to suggest that not all causal explanations give reasons.

One way in which reason-giving explanations differ from other causal explanations is that they make sense of the action from the agent's point of view.¹³ This idea of making sense of or rationalizing does not seem to have any application outside of someone's perspective. There's what it makes sense to expect given that you know about certain causal conditions, e.g., the faulty valve. But those conditions by themselves do not rationalize or make sense of the explosion. So one side to the notion of a reason is the explanatory side. Reasons figure in the kind of explanations that explain not just by pointing to causal conditions, but by pointing to causal conditions that rationalize or make sense of what happens.

The other side to the notion of a reason is the normative side. We can often say that people have reasons, or good reasons, to do something we know they're not going to do. Often, our concern with attributing reasons to people is much more concerned with what they ought to do than it is with what they're going to do. This is the normative side to the notion of a reason. It's just a sad fact of life that people sometimes act (and believe) for bad reasons. In such a case there will be a reason-giving explanation, but we won't want to say anything particularly positive in the normative dimension. And it's another sad fact that people don't always do (or believe) what they should, or what they have most good reason to do (or believe). Here there's no explanation of their doing it, but there is something to be said in favor of their doing it.

¹³ Davidson (1980).

The sad facts of life require us to distinguish good reasons from bad reasons. But they don't require us to think that good and bad reasons are of distinct ontological categories, e.g., that good reasons are beliefs and bad reasons are facts or the other way around. Perhaps if no coherent conception of a reason, consistent with the facts, could unify the two sides to the notion, we might have to reluctantly conclude that there are these two quite different kinds of reasons. Motivating reasons explain but do not confer positive normative status. And normative reasons confer the status but do not explain.

Before we adopt this view, however, we should at least look to see if there is a coherent, unified account of reasons consistent with the facts. It turns out there is. My basic idea of a reason is relatively straightforward. Reasons are supposed to make things reasonable. Good reasons do what they're supposed to do. They're things that make things reasonable. Bad reasons are still reasons: they're supposed to make things reasonable. But they don't do what they're supposed to. So suppose you believe that p and that if p then q . And suppose this is what makes you believe that q . What your belief is based on is supposed to make it reasonable. If either of the earlier beliefs is unreasonable or unjustified then they won't make your belief that q reasonable. But if the beliefs themselves are reasonable, and if the transition from them to the conclusion is reasonable, then your belief that q will also be reasonable, and it will be reasonable in virtue of where it came from.

Suppose you want E and think that M -ing will get you E , and this belief-desire pair makes you M . What makes you do it is supposed to make it reasonable. If there's something unreasonable about either of the earlier mental states, then there's something unreasonable about M -ing as well. But if the earlier mental states are reasonable, and if, in the context of your current perspective on things, the transition from them to M -ing is reasonable as well, then M -ing is a reasonable thing to do in those circumstances. Maybe this transition is not reasonable. Maybe the mere desire for an end doesn't make it reasonable to pursue it. Maybe you need the belief that the end will be good, or maybe you just need the belief that you will want the end when you get it. Maybe desires themselves can be unreasonable, and maybe they can't. These are all perfectly good questions. And this is precisely what the study of practical reason ought to be: the study of which mental states make it reasonable for you to act.

Even given this brief sketch, it's clear that on this picture, the very same reasons have both an explanatory and normative side. What makes you do it or believe it determines whether the action or belief is reasonable. If being reasonable is not a positive normative status, I don't know what is. In one sense the two sides can come apart. If

you act (or believe) for bad reasons, then what makes you do it fails to make it reasonable. And if you don't do what you have good reason to do, then the mental states that make it reasonable for you to act (or believe) fail to make you do it. But that doesn't mean that good and bad reasons are different kinds of things. So the sad facts of life by themselves can't motivate dualism about reasons, and neither can the absence of an alternative, unified account. What motivates dualism about reasons is the idea that, for example, in the dairy thief case where you mistakenly think you have milk, there must be something to be said for going to the store for milk. And since this something can be as inaccessible to the agent as you like, it cannot be part of the agent's perspective, and so, at least according to *The Bad*, it cannot be the kind of thing that makes things reasonable.

There are two main questions to ask about this conception of reasons, the conception of reasons according to which motivating reasons are perspective-dependent and determine the rationality of action while normative reasons are perspective-independent and let's just say, do something else. If normative reasons don't determine the rationality of action, what positive normative status do they confer? And what should we do when the rationality of action conflicts with the mystery status? These are my questions.

What Good Are They?

Since the defining feature of *The Bad* is the gap between normative reasons and rationality – the former are perspective-independent while the latter is perspective-dependent – whatever positive normative status is conferred by normative reasons, that status cannot be rationality. It's just part of the view that in the dairy thief case, you have a normative reason to go to the store even though going to the store would be irrational. But there is something to be said for going to the store in that case. And the question of what there is to be said is the question of what positive normative status is conferred by normative reasons.

I only know of two possibilities for the relevant positive status. I have no argument to show that it could only be one of these two possibilities. These are just the only ones I've ever heard of or can think of. Maybe the relevant positive status is evaluative and not really normative at all. If you went to the store in the dairy thief case and bought some milk by mistake, this would not be a reasonable thing to do. But it would be a good thing. So according to what we might call the value-based conception of reasons, to say that you have a normative reason to A is basically to say that it would be a good thing if you

A-ed. It can be a good thing for you to A even when you think that A-ing would be disastrous.¹⁴

If you went to the store for milk in the dairy thief case, this would not be a reasonable thing to do. But we can say this much for going to the store. If you knew all the facts in the dairy thief case, then going to the store would be a reasonable thing to do. According to what we might as well call the omniscience-based conception of reasons, to say that you have a normative reason to A is basically to say that A-ing would be reasonable if you knew all the facts.¹⁵

The problem for both of these views is the same. Each has trouble explaining the simple idea that the notion of a reason is a normative notion. This is easiest to see in the case of the value-based conception, but it applies to the omniscience-based conception as well. If your conception of something can't make sense of it's being normative, then it's not a conception of a reason, and no amount of terminological stipulation will make it one.

Value

I have no idea whether or not there might someday be a complicated philosophical account of normative notions like right, ought, should, reason, and reasonable in terms of evaluative notions like good, better, and best. But I do know enough to know that the Simple Identification must be false.

(SI) S ought to ϕ iff S's ϕ -ing would be best.

There's an avalanche coming down the mountain, about to hit a ridge. If it goes to the left, it will destroy the village. If it goes to the right, it won't. Clearly it would be better if the avalanche went to the right. Should we say that the avalanche ought to go to the right or has a reason to go to the right? Or should we just say that there is a reason for the avalanche to go to the right, but the avalanche, lacking a mind, just can't see it? If we shouldn't say these things then (SI) is false and we need to add some restriction on the implicit quantifiers binding "S" and " ϕ ." But why shouldn't we say these things and what justifies the restriction on the quantifiers?

The value of the destruction of the village depends in obvious ways on mental states. At least part of what makes it a bad thing is all the pain and suffering involved. But evaluative notions like good and bad

¹⁴ Both Dancy (2000) and Raz (2005) seem to accept some form of the value-based conception of reasons.

¹⁵ Smith (1994) seems to accept the omniscience-based conception of reasons.

apply in the first instance to states of affairs, and whether there are mental states in the causal history of the state of affairs seems to be largely irrelevant. The villagers' mental states might matter, but the avalanche's mental states, or lack thereof, doesn't. Of course, the consequences of a state of affairs might affect its value. So evaluative notions may well look forward, but they do not appear to look backward.

Normative notions, by contrast, apply in the first instance to people, agents, believers, or things with minds. It's just not true that the avalanche ought to see what will happen if it goes to the left, and it's not true that the avalanche ought to go to the right. But the idea that only things with minds have reasons can't just be an afterthought. It can't just be something you tack onto the end of your theory to avoid a counterexample. If your basic idea of a reason doesn't explain the connection between reasons and minds, it's time to get a new idea.

Can the value-based theorist explain the connection between reasons and minds? Well, the theorist can assert that only things with minds have reasons. Simply restrict the range of the variable "S" in (SI) to things with minds and you avoid the counterexample. But can the theorist explain the restriction? If I say that only people with banjos have reasons, but my account of what reasons people have is completely independent of the state of their banjo, you should be deeply suspicious of my adopting a restriction that looks completely ad hoc from the point of view of my own theory.

If you march into your local nuclear power plant and press the shiny red button that is clearly labeled "Self-Destruct Mechanism: Do Not Touch," then on this one in a million occasion, instead of destroying the plant, it will keep the avalanche from destroying the village. You, of course, have no way of knowing this; you reasonably believe that the consequences of pressing the button will be disastrous; and you have no interest in and see nothing good about killing yourself or others. To the extent that your mental states make a difference to your reasons, they all argue against pressing the button. But since your pressing the button would be a good thing, the value-based conception of reasons entails that you do have a reason to press the button. You don't have one of those silly old motivating reasons. You have a *normative* reason. Since knowledge, evidence, and justified belief have nothing to do with normative reasons, we can easily tell the story so that you have no better reason not to press the button. So in this case, the account entails that you ought to press the button.

Suppose you do what they say that you should. You march zombie-like into the plant and press the button. Given your lack of motivating reasons, this might be an elaborate nervous twitch, but it couldn't

possibly be an intentional action. So whatever these reasons are, they don't look like reasons for action. They may be reasons for your body to be in a certain place at a certain time, or they may be reasons for certain states of affairs to obtain. But whether those states of affairs obtain (or whether your body ends up where it is) as a result of an action, intentional or otherwise, is completely irrelevant. If you can comply with the reason without believing, then it cannot be a reason to believe. If you can comply with the reason without acting, it cannot be a reason for action. If they're not reasons for acting, and they're not reasons for wanting or intending to act, it's hard to see how these so called reasons are practical.

If the theory says that you have a reason to press the button in virtue of the fact that pressing would be a good thing, then the theory is committed to treating like cases alike. This is part of what "in virtue of" means. If the avalanche's going right of its own accord would be a good thing then the avalanche has a reason to go to the right, and the theory cannot restrict "S" to agents. If the theory says instead that you have a reason to press the button partly in virtue of the fact that pressing the button would be a good thing and partly in virtue of the fact that you have a mind, then the presence of a mind must make a difference to the presence of a reason. This is also part of what "in virtue of" means. But if, from the point of view of the theory, it makes no difference whether you press the button on purpose, whether your pressing the button is an action, or what mental states you happen to have, then it is, to say the least, completely unclear how the theory could explain how the presence of a mind could make a difference to the presence of a reason. So either avalanches have the same sorts of reasons that you do, or the value-based conception of reasons must go.

Omniscience

Suppose you find out that there's a relatively easy way to end world hunger but you have to go all the way to the library to find out what it is. Do you have a reason to go to the library? We can assume, if you like, that you want to end world hunger; that there's nothing wrong with ending world hunger; and that you're not currently doing anything more important than ending world hunger. Do you have a reason to go to the library yet?

If claims about what you have reason to do are about what would be reasonable for you to do if you knew all the facts, then you have no reason to go to the library. If you knew all the relevant facts, you wouldn't need to go to the library to find out how to end world hunger. You would already know. If this is a problem, it doesn't merely

look like a problem with some specific formulation of the omniscience-based account of reasons. It looks like a problem for the basic idea.

Suppose that instead of asking what would be reasonable for you to do if you knew all the facts, we adopt Smith's version of the view.¹⁶ Now we ask what you, if you knew all the facts and were perfectly reasonable, would advise your lesser self to do in your lesser self's actual circumstances. If I were perfectly reasonable and knew all the facts, I would tell myself how to end world hunger, and that's what I would want myself to do. I certainly wouldn't tell myself or want myself to waste a half an hour going to the library to find the answer to a question that I could just as easily answer immediately. Given the facts about world hunger, a disturbing number of people can die in a half an hour. I would never tell myself to do something simply in order to find something out. I would tell myself what I wanted to know. I would, however, tell myself to press the self-destruct button in the local nuclear power plant if that would avert the destruction of the village or end world hunger.

So if the idea is that we ask what I, if I knew all the facts, would advise my lesser self to do, then it's fairly clear that we never have any reason to learn anything. But suppose we ask what I, if I knew all the facts, would *want* my lesser self to do. Wouldn't you want your lesser self to go to the library? Of course you would. So maybe this version of the view gets out of the problem.¹⁷ But wouldn't you also want your lesser self to press the shiny red button if that would end world hunger? Of course you would. And given the time constraints, which would you want more? Presumably, you'd want yourself to press the button, just as you're supposed to want your lesser self to buy the expensive painting your lesser self has no interest in. The fact that it's not going to happen is not something you're supposed to take into account when you know all the facts, and neither is the fact that it would be irrational if it did. So perhaps this version of the view can deliver the idea that you have *some* reason to find things out. But it doesn't look like it can deliver the idea that you ever have *most* reason to find things out. Doing the implausible, irrational thing will always produce better results, and all-knowing you will know that. So it will never be true that you ought to go to the library, or that you ought to ask directions, or that you ought to learn anything at all.

I don't know who would take this as an objection. I find it no more surprising to be told that I never have most reason to find

¹⁶ Smith (1994): Ch. 5.

¹⁷ I'd like to thank an anonymous referee for *Philosophy and Phenomenological Research* for raising this objection.

anything out than I find it to be told that I do have a reason to press the self-destruct button. So I don't know what it would be like to find one of these consequences intuitive and the other counter-intuitive. But suppose you do. In order to fix your theory to avoid the counter-intuitive consequence, you need some clue, some very rough idea of what business reasons are in. They're not in the business of making things reasonable. They're not in the business of making things good. And they're not in the business of making things reasonable-if-you-knew-all-the-facts. Can you really be sure that there's anything left to have in mind?

So suppose that you don't take it as an objection. Or suppose that you do, but someday, someone somewhere comes up with a minor revision of the theory that allows you to have reasons to find things out but still takes as its basic idea the idea that reasons make things reasonable-if-you-knew-all-the-facts. Now suppose that I say of some particular proposition that it would be known by someone who knew all the facts. Clearly, I use epistemic words, "known" and "knew" in describing the proposition. But have I attributed any genuinely epistemic property to this proposition? Have I really said anything more than that it's true? We may use a normative word when we say that pressing the button would be reasonable if you knew all the facts. But have we really attributed a normative property?

Imagine what it would be like to do what the omniscience-based theorist says that you ought to do. You march into the plant and press the button. We're not to imagine you in a world where you know all the facts. We're not even to imagine you in a world where you've just had a chat with an all-knowing version of yourself. We're to imagine you in a world, as close to the actual world as possible where it's true that, though you have no way of knowing that, if you knew all the facts you would do it or want yourself to do it. Given the lack of motivating reasons, you couldn't press the button on purpose. But so what? If you press the button by mistake, you've done what all-knowing you wants you to do. And if someone shoves you, and you push the button with your face while slamming into the wall, you've done what all-knowing you wants you to do. How could this theory care about whether your pressing the button was the result of an action or not?

Of course the theory cares about your mental states in the far off world where you're perfectly reasonable and know all the facts. But how could the theory care about your mental states in the closer world in which you have a reason to press the button but in which pressing the button could at best be a mistake? And if you can satisfy all-knowing you without acting, how could it possibly matter whether you're asleep, in a coma, or recently deceased? Again, if you can

comply with it without acting, then it can't be a reason for acting. And if the presence or absence of a mind is irrelevant to your having it, then it can't be a reason. We can use a normative word to talk about these things. They're things that make things reasonable-if-you-knew-all-the-facts. But that doesn't make them normative things.

Conflict

If we try to give an account of the normative status conferred by normative reasons either in terms of value or in terms of omniscience, it will turn out that normative reasons are neither normative nor reasons. Perhaps there's a third alternative. Whatever the alternative, the problem will be the same. Imagine a case where they say you have a reason but have no way of knowing that you have a reason. Now imagine doing what they tell you to do. Since you can't do it on purpose, it's hard to see how it could matter if it's an action. And if it doesn't matter if it's an action, it's hard to see how it could matter if you have a mind. But if none of these things matter, it's hard to see how it could be a normative reason for action.

But even without knowing very much about the positive normative status allegedly conferred by normative reasons, we can ask about what happens when that status conflicts with another positive normative status: being reasonable. We can say that an action is *wonderful* when it has the status conferred by normative reasons, and we can say this without worrying too much about what wonderfulness comes to. What happens when an action is wonderful but not reasonable or reasonable but not wonderful? What should you do?

Consider the following pair of ideas.

- (A) If you have most (good) reason to ϕ , you ought to ϕ .
- (B) If it's unreasonable for you to ϕ , you shouldn't ϕ .

At the very least, I think we should prefer a view consistent with both of these ideas to a view that forces us to choose between them. If good reasons are things that make things reasonable, then (A) and (B) go together nicely. But if normative reasons make things wonderful, and pressing the self-destruct button in the nuclear power plant is wonderful but unreasonable, then either we go with (A) and say that you ought to press the button, or we go with (B) and say that you shouldn't.

If you're thinking about giving up (B), read it again. Sometimes it seems as though some people (usually young non-philosophers) think

that Reason and Emotion are enemies. There's just no room in their minds for the thought that loving people can give you a reason to be with them. Some of these people side with Emotion over Reason in the great battle between the two. If you tell them they're being unreasonable about something, they agree with you and say, "So what?" I don't think this theory has much going for it, but at least I understand it. Now imagine a group of analytic philosophers presenting arguments, evaluating counterexamples, and devising theories about the nature of practical reason. They end up rejecting (B). You tell them they're being unreasonable about something. They agree with you and say, "So what? What does the fact that it's unreasonable for me to ϕ have to do with what I ought to do?" In this case, it's much more difficult to figure out what could possibly be going on.

"Be reasonable" is a categorical imperative. It applies to all agents regardless of their contingent desires. If a code of conduct says that you ought to ϕ in circumstances in which it would be obviously and completely irrational for you to ϕ , then you are rationally required to ignore that code of conduct, regardless of whether it comes from your parents, your favorite moral theory, or the voices in your head claiming to be from God. If you are rationally required to be reasonable, and if rational requirements are genuine requirements (even if they're not the only genuine requirements) then denying (B) is out of the question.

As far as I can tell, it makes no difference how many ways ambiguous the English word "ought" may be. Suppose, for the sake of argument, that there's a sense of "ought" in which the claim that you ought to ϕ means that the rules of etiquette require you to ϕ . And suppose that you discover that in many circumstances you ought, in that that sense, to ϕ even though ϕ -ing in those circumstances would be clearly and obviously irrational. Now ask the following two questions. What is the only rationally permissible response to this discovery? And how do you in fact respond to this discovery? The actual response is the rationally required response. The discovery is treated as a decisive reason to doubt the normative significance of both the relevant sense of "ought" and the code of conduct. If there is an objective sense of "ought" defined in terms of normative reasons, then the corresponding discovery ought (in a sense of the word you care about whether you want to or not) to lead to the corresponding response.

If your conception of reasons requires you to give up one of the pair, but accepting (B) is rationally required, then the only possible alternative left is to reject (A). So you've got these things, normative reasons, whose status both as normative and as reasons is already in doubt. Now you're willing to admit that they don't tell you what to do

or make it the case that you ought to do one thing rather than another. If you're willing to admit that, you should be willing to admit that they're not practical reasons at all.

If you can't choose between (A) and (B), but can't stand the idea that practical reasons could be determined by your perspective, you might try out the idea that marching into the plant and pressing the self-destruct button is a reasonable thing to do, even though you have no way of knowing that it's a reasonable thing to do. While you're at it, you might try out the idea that believing there's an even number of blades of grass in Washington Square Park is perfectly reasonable if you happen to get it right, even though you have no way of knowing whether it's reasonable or not. This involves giving up *The Bad* and turning to *The Ugly*. But it also involves giving up on the ideas of justification and rationality. Guessing correctly is not a way of being justified, and truth is not a kind of justification. Truth is what skeptics have instead of justification. If you don't believe in rationality, that's your business. But if you really don't believe in rationality, it's unclear what you're talking about when you go on about practical reason.

Either normative reasons don't tell you what to do at all, or they tell you the wrong thing. The only question left to ask about them is how anyone could have believed in them in the first place. There are various arguments in the literature against the idea that practical reasons are psychological states. To the extent that *The Good* is committed to psychologism, these are indirectly arguments against *The Good*. Let's take a look at these arguments.

The Phenomenology of Deliberation

Usually when we're deliberating about what to do, we're thinking about the world. We're not thinking about our own mental states. But this isn't always true. In some cases, let's call them the unusual cases, you take the fact that you're in a certain mental state as a reason to do something. For example, you might take the fact that you believe that everyone is out to get you as a reason to see a psychologist. Presumably, if you took the fact that everyone is out to get you as a reason to do anything, you'd take it as a reason to go into hiding. That would be thinking about the world, and it would be a usual case. But if you take the fact that you believe they're out to get you as a reason to seek help, then you're thinking about your own mind, and we have an unusual case.

But if psychologism is the view that reasons are psychological states, doesn't it follow that all cases are unusual cases? At the very least, the psychologist must give some kind of account of the difference between

the usual and unusual cases.¹⁸ Though there's very little agreement about the notion of a reason, I take it that the notion of reasoning or deliberating is relatively uncontested. Reasoning is a psychological process. It goes on in the mind. And it involves transitions from some mental states to others. The question is not about what reasoning is. The question is about how reasons figure in reasoning.

Perhaps the objection relies on an implicit argument to the effect that psychologism entails that all reasoning involves mental states that are about other mental states, i.e., that all reasoning is from second-order mental states. Let's see if we can make this argument explicit. Here's one attempt.

- (1a) For a reason to figure in reasoning is for it to be the content of one of the mental states involved in the psychological process of reasoning.
 - (1b) Psychologism entails that your belief that p (or the fact that you believe that p) is a reason for you to ϕ .
-
- (1c) So psychologism entails that your belief that p can figure in reasoning only if it is the content of a second-order belief.

It's useful to compare that argument with the following argument designed to show that objectivism entails anti-realism.

- (2a) For a reason to figure in reasoning is for it to *be* one of the mental states involved in the psychological process of reasoning.
 - (2b) Objectivism entails that the obviously external fact that you're out of milk is a reason for you to go to the store.
-
- (2c) So objectivism entails that obviously external facts are psychological states.

I take it that no one takes the second argument seriously. The first premise simply assumes an account of what it is for a reason to figure in reasoning that the argument's opponent would not and could not accept. But at least in this respect, there seems to be no difference between the first argument and the second.

¹⁸ Versions of this argument can be found in Collins (1997), Dancy (2000), Hyman (1999), Scanlon (1998), and Thomson (2003).

Perhaps there's some independent reason for adopting one of these conceptions of what it is for a reason to figure in reasoning. Indeed there is. Consider the following two cases. Case 1: You believe that p ; you believe that if p then q ; and for these reasons, you conclude that q . Case 2: You believe that p ; you merely hope that if p then q ; and for these reasons you conclude that q . If reasons are contents, the reasons are the same because the contents are the same. If they're good reasons in the first case, they must be good reasons in the second case as well. But we know they're not good reasons in the second case. And we know that what's wrong with the reasoning has nothing to do with the contents. The problem is with the attitudes you take toward those contents. So no account of reasons that refuses to talk about the nature of the attitudes can distinguish good reasons from bad.

Perhaps I have not accurately represented the intended argument. But without some account of what it is for a reason to figure in reasoning, an account that is at least quite likely to be inconsistent on its own with psychologism, it's very hard to see how you could get from something like (1b) to something like (1c). Perhaps no argument was intended. Perhaps it was simply an invitation for the psychologist to explain the difference between the usual and unusual cases. This is a reasonable request. The difference is that in the usual case your reasoning involves first-order propositional attitudes, and those are your reasons. In the unusual case, your reasoning involves second-order propositional attitudes, and those, the second-order propositional attitudes, not the first-order attitudes they're about, are your reasons. That's the difference.

Advice

Suppose I come to you for advice. In telling me what to do or what I ought to do, you need not restrict your attention to what I'm already in a position to know. If you know that M is a means to my end, then you can tell me that I ought to M even if you're quite sure that I had no way of knowing that M would get me what I want. Doesn't this show that what I ought to do, and by extension, what I have reason to do is not determined by my perspective?¹⁹

The first thing to notice about this argument is that it assumes a close connection between reasons and oughts. To the extent that this argument is in the service of The Bad, the closer the connection between reasons and oughts, the looser the connection between what's reasonable and what you ought to do. If accepting this argument

¹⁹ See, for example, Thomson (2003).

requires rejecting the idea that if you ought to ϕ , it's at least reasonable for you to ϕ , then you should be somewhat suspicious. In order for the argument to work, your telling me that I ought to M can't make a difference to what I ought to do. It must have been true all along that I ought to M, even though before you gave me the advice, doing that would have been completely unreasonable.

There's no doubt that we use "ought" to give advice without restricting attention to the perspective-dependent facts. The question is what we should conclude from this fact. We can give advice without using the word "ought." At least everywhere I've ever lived, the standard way to give someone directions is by saying things like, "You want to turn left at the light." Using "want" to give advice is not restricted to driving. "You want to show up early because they start getting busy around 7." Since we can use "want" to give this advice without restricting our attention to the perspective-dependent facts, should we conclude that desires are not and are not determined by your mental states? We certainly shouldn't conclude this if there's a reasonable alternative.

In any ordinary case, if you give me directions by telling me that I want to turn left at the light, then at the instant you utter the sentence, what you say is not strictly and literally true. If I already wanted to go left at the light, I wouldn't need directions from you. But a second after you utter the sentence, when I understand it, then what the sentence says is true. If I'm following your directions, then once I understand them, I do want to turn left at the light. I don't think that the one-second gap between utterance and fact is a serious reason for people to stop speaking this way. And I don't think that the fact that people speak this way is a serious reason to doubt the mentality of desire.

Perhaps something similar goes on with "ought." Before you told me that I ought to M, it would have been completely unreasonable for me to M. So it wasn't strictly and literally true that I ought to M. But once I understand what you're saying and why you're saying it, you've changed what facts are accessible or available to me. As a result, you've changed what reasons I have. So it looks as though there are two options. According to The Bad, I ought to M even when it's completely unreasonable for me to M. According to the alternative, giving people advice changes their practical situation. The alternative has the added benefit of explaining why people ask for advice in the first place. Of the two options, I take it that only the alternative is rationally permissible. If you're thinking about giving up on rationality, this consideration might not be as compelling as it seems. But if you really were thinking about giving up on rationality, you probably would not be giving arguments in favor of your view.

The Good

If your reasons are determined by the facts while what's reasonable is determined by your beliefs, then you will often have reason to do completely unreasonable things. If what you ought to do is determined by what you have best reason to do, then the theory will tell you that you ought to ϕ in circumstances in which ϕ -ing would be completely and obviously irrational. If "Be reasonable" is a categorical imperative, you are rationally prohibited from doing what they tell you to do. All of this is a result of trying to drive a wedge between reasons and what's reasonable. So The Bad and dualism about practical reasons must go.

Can we hold onto The Good and still identify reasons with perspective-independent facts? It doesn't look that way. Take the case of a justified, false belief. If it's reasonable for you to believe that you're out of milk then it's reasonable for you to act on that belief. So suppose that you go to the store in the dairy anti-thief case. If reasons are facts, it looks as though there is no reason for you to go to the store for milk. So it looks as though you have no reason to go to the store. So you don't go to the store for a reason. So you don't go to the store intentionally. But this is obviously false. You might try to open up a gap between doing something intentionally and doing it for a reason. But if you say that there are these mental states that make things intelligible, make things intentional, and make things reasonable, but you would prefer not to call them reasons, then your view is a notational variant of dualism. You believe in motivating reasons, you just don't call them that.

Jonathan Dancy has a different account of the case of false belief.²⁰ Since he is unwilling to identify reasons with psychological states, we're told that sentences like, "She went to the store because she was out of milk" and "Her reason for going to the store was that she was out of milk" can be true even when she was not out of milk.²¹ Some explanations are factive, and some are not. Some explain in terms of things that happen, and some explain in terms of things that don't happen. In addition to the inherent implausibility of the view, it looks like a rejection of the idea that reasons are facts. Whatever "that she was out of milk" refers to in these circumstances, it's not a fact, state of affairs, or objective feature of the circumstances.

So what's reasonable for you to do and what reasons you have are both determined by your perspective. And reasons are either psychological states, or facts about psychological states, or the contents of psychological states. But which psychological states are reasons and

²⁰ Dancy (2000): Ch. 6.

²¹ Ibid p. 132

which ones are not? This is an enormous question, and answering it involves figuring out the role of desire and normative belief in the determination of practical reason. Pretty much all of this will have to wait for another occasion, but I do think it's important to include knowledge, ordinary knowledge of ordinary facts, in our list of psychological states that constitute reasons. There is a familiar picture of the mind as a self-contained, inner realm whose nature is completely independent of what goes on in the external world. There are various reasons for rejecting this picture.²² And rejecting it here is just as important as rejecting it everywhere else.

I'm primarily concerned with defending *The Good*, where this is understood as a claim about what's made true by your perspective. The basic idea behind a perspective or point of view is the idea of what you can see from where you are. What you can see is not determined by what you're like on the inside, in any sense of that expression. It's not determined by your intrinsic properties, and it's not determined by what's available to you on the basis of introspection and a priori reasoning. It does partly depend on what you're like on the inside, at least how well your eyes are functioning. But it also depends on what there is in the world to see.

The idea of what you can see from where you are generalizes in certain obvious ways: what you can perceive from where you are, what you can know from where you are, or what you're in a position to know, and what you're in a position to do on purpose. This last, like the others, depends not only on what you're like on the inside, or what the world is like on the outside, but on the relations between the inner and the outer that make minds worth having. Any attempt to understand your perspective in terms of your non-factive mental states would simply be a rejection of the basic idea.

So it might not matter that much whether we call these perspectival facts mental or psychological facts. In this context, probably the best reason to think of knowledge as a mental state is that it provides unity to our account of practical reasons. We'd like an account of the following two sorts of claims, not only an account that allowed them both to be true, at least occasionally, under ordinary circumstances, but also an account that told us what they have to do with each other.

(S) She went to the store because she thought she was out of milk.

(O) She went to the store because she was out of milk.

²² See, for example, Putnam (1975), and Burge (1977).

In order for (S) to be true, she must believe that she's out of milk, and this belief must cause her to go to the store. But not just any old kind of causation will do. Her belief must be causally related in the appropriate way to the action. Of course, it would be nice to have an account of which way is appropriate. But at least we have a basic idea of what makes (S) true: the mental state causes, explains, and rationalizes the action.

If (O) is an ordinary, reason-giving explanation of intentional action, then in order for (O) to be true, she must know that she's out of milk. If you act in light of the fact that p then you must be aware that p . And if you're aware that p , you must know that p .²³ Of course that knowledge must be causally related in the appropriate way to her action. But a unified account of (S) and (O) is available. In either case, the relevant mental state, belief in the former and knowledge in the latter, must cause, explain, and rationalize the action.

Wait a minute. If the fact explains the knowledge, and the knowledge explains the action, doesn't the fact explain the action as well? If both the fact and the knowledge explain the action, why identify the reason with one rather than the other? In this particular case, both the fact and the knowledge explain the action. But considerations of generality suggest that knowledge is the real causal power at work. Suppose she goes to the store now because she'll be out of milk by tomorrow. The future fact doesn't cause the present knowledge or action. And while I don't think that causal explanation is the only kind of explanation, I can't quite see how the future fact could explain the present knowledge or action. Similarly, if she says "twelve" in response to his question because $5 + 7 = 12$, the mathematical fact doesn't cause, and so it's hard to see how it explains, the knowledge or action. Knowledge explains in all of these cases while the facts do not, and all of these cases ought to be treated alike. And of course, when you have the fact without any awareness of the fact, the fact doesn't explain the action at all. So knowledge provides a better explanation than the facts even in those cases where the fact is part of the causal history of the action.

So you could agree with me on the conditions under which someone has a reason and on what it takes for someone to act for a reason but continue to insist that reasons must be facts. When she knows she's out of milk, her reason is the fact that she's out of milk. When she falsely believes she's out of milk, you can't identify her reason with the fact that she's out of milk, since there is no such fact, and you can't identify it with the false proposition that she's out of milk, since that's not a fact. But you might try identifying her reason not with her belief

²³ This argument is in Hyman (1999).

that she's out of milk, but with the fact that she believes that she's out of milk.²⁴ This requires arguing not only that there is a difference between your belief that *p* and the fact that you believe that *p*, but that this difference actually matters in the present context.

But if you agree with me on all the substantive issues, e.g., that you can't have most (good) reason to ϕ in circumstances in which ϕ -ing would be completely and obviously irrational and that what's reasonable for you to do in your actual circumstances is determined by your actual perspective rather than the perspective of someone in completely different epistemic and practical circumstances, then the only thing left to fight about is who gets the word "reason." Here I'm willing to be generous. If you can't call the relevant mental states reasons, call them things that make things reasonable. If you can't call knowledge a mental state, call it an element of your perspective. Either way, your view will be a notational variant of mine.²⁵

References

- Burge, Tyler (1977), "Individualism and the Mental," *Midwest Studies in Philosophy* vol. 4, eds. French et al, (Minneapolis: Minnesota University Press): 73–121.
- Collins, Arthur (1997) "The Psychological Reality of Reasons," *Ratio* 10: 108–123.
- Dancy, Jonathan (2000), *Practical Reality* (Oxford: Oxford University Press).
- Davidson, Donald (1980), "Actions, Reasons, and Causes," in *Essays on Actions and Events* (Oxford: Clarendon).
- Gibbons, John (2001), "Knowledge in Action" *Philosophy and Phenomenological Research* LXII.
- Hyman, John (1999), "How Knowledge Works," *Philosophical Quarterly* 49: 433–451.
- Korsgaard, Christine (1986), "Skepticism About Practical Reason," *Journal of Philosophy* 83: 5–25.
- McDowell, John (1995), "Knowledge and the Internal" *Philosophy and Phenomenological Research* 55: 877–893.
- Parfit, Derek (1997) "Reasons and Motivation," *Proceedings of the Aristotelian Society Supplementary Volume* 71: 99–130.
- (ms.) *Climbing The Mountain*

²⁴ Dancy actually toys with this idea in various places in his (2000).

²⁵ I'd like to thank Joe Mendola, Mark van Roojen, the members of my Practical Reasons seminar, and an anonymous referee for *Philosophy and Phenomenological Research* for helpful comments and suggestions.

- Putnam, Hilary (1975), "The Meaning of 'Meaning,'" in *Mind, language, and Reality*. (Cambridge: Cambridge University Press).
- Raz, Joseph (2005), "The Myth of Instrumental Rationality." *Journal of Ethics and Social Philosophy* 1: 2–28.
- Scanlon, T.M. (1998), *What We Owe to Each Other* (Cambridge Mass.: Harvard University Press).
- Smith, Michael (1994), *The Moral Problem* (Oxford: Blackwell).
- Thomson, Judith (2003), *Goodness and Advice* (Princeton: Princeton University Press).
- Wiland, Eric (2003), "Psychologism, Practical Reason, and the Possibility of Error," *Philosophical Quarterly* 53: 68–78.
- Williams, Bernard (1980), "Internal and External Reasons," in *Moral Luck* (Cambridge: Cambridge University Press).
- (1989) "Internal Reasons and the Obscurity of Blame," in *Making Sense of Humanity* (Cambridge: Cambridge University Press): 35–45.
- Williamson, Timothy (1995) "Is Knowing a State of Mind?" *Mind* 104: 533–565.