

THE SCOPE OF INSTRUMENTAL REASON

Mark Schroeder
Princeton University

Allow me to rehearse a familiar scenario. We all know that which ends you have has something to do with what you ought to do. If Ronnie is keen on dancing but Bradley can't stand it, then the fact that there will be dancing at the party tonight affects what Ronnie and Bradley ought to do in different ways. In short, **(HI)** you ought, if you have the end, to take the means. But now trouble looms: what if you have dreadful, murderous ends? Ought you to take the means to them? Seemingly not. But fortunately, an assumption made by deontic logics¹ comes to the rescue. Since "ought", according to this assumption, is a sentential operator, HI must really be *ambiguous*. It could be read either as **(Narrow)** You have the end \rightarrow O(you take the means) or as **(Wide)** O(you have the end \rightarrow you take the means). Now if Narrow is true, then you really ought to take the means to your murderous ends. But this doesn't follow from Wide. All that follows from Wide is that you ought to either take the means to these ends or else give them up. Conclusions: (1) Since HI is on some reading true, but Narrow isn't, Wide is true. (2) Wide accounts for the relationship between your ends and what you ought to do.

This elegant scenario repeats itself in many other domains in which it seems like something can have a *bearing* on what some particular agent ought to do. Does what you know affect what you ought to do? Do your beliefs about what you ought to do affect what you ought to do? Do your promises affect what you ought to do? Do your beliefs affect what you ought to believe? On each of these counts, the intuitive answer is "yes". And so each of these questions leaves something for the moral philosopher or the epistemologist to investigate. On each count, it seems that what we all know, is that **(Account)** you ought, if p , to do A . But on each count, the Narrow-scope reading of the "ought" in this claim yields unintuitive consequences. So since Account is true, it must be true on the Wide-scope reading. So the Wide-scope principle must be what accounts for what each of these things has to do with what you ought to do. The Wide-Scope program in the theory of practical and theoretical rationality is to offer these

kinds of account, in answer to questions of the form, “what does this or that have to do with what you ought to do?”

Proponents of Wide-Scoping² hold that this motivation is conclusive. Indeed, they say that Wide-scope principles are “uncontroversial”³ and that those who do not agree are “confused”. It would be nice if this were so. But unfortunately the issues are in fact somewhat more complicated. In this paper, I will explain why the Wide-Scoping program is highly controversial. Just as the Narrow principles yield unintuitive results, their Wide counterparts invoke two different kinds of highly controversial commitment. Indeed, I’ll argue, on an acceptable semantics for “ought”, HI and its counterparts in the other domains are *not*, after all, ambiguous between Narrow and Wide readings. This means that if the Narrow principles are false, so are HI and its counterparts. If this is right, then the Wide-Scope views aren’t so much offering an *interpretation* of the uncontroversial HI, but are rather *replacing* HI with an alternative that is weaker in one respect, in the face of counterexamples.

Once we see that the project is that of *weakening* HI, however, rather than that of *interpreting* it, I’ll suggest that this approach is narrow-minded. If we take a more broad-minded approach, we can see that another way of weakening HI is possible. Instead of weakening it by looking for wider *scopes* for the “ought”, we can weaken it by replacing the “ought” itself with a weaker normative concept, such as that of a *reason*. I’ll suggest that this is a perfectly viable kind of response to the counterexamples to HI and its analogues. Since this kind of response also escapes the highly controversial features of the Wide-Scope accounts, I hold that these kinds of account are actually to be *preferred*. And this, as it turns out, has implications for at least one fundamental and hotly-debated issue in moral philosophy.

1. Objective Instrumental Rationality

When Ronnie is keen on dancing and Bradley can’t stand it, the fact that there will be dancing at the party affects what each ought to do differently. So being keen on dancing must somehow relate to what one ought to do. But how? That is the question to be answered by an account of objective instrumental rationality. It is an account of instrumental rationality, because it tells us what your ends or desires have to do with what you ought to do. It is an account of *objective* instrumental rationality, because it has to do with what actions are *actually* means to your ends, rather than with what actions you merely *believe* to be means to your ends. That question is answered by an account of *subjective* instrumental rationality, and we’ll return to it in section 2.

According to a naïve view, we *already* know at least a *little* bit about how to account for the domain of objective instrumental rationality. It is that HI is true. You ought, if you have the end, to take the means. As noted, if we assume that “ought” takes propositions for one of its relata, then this can receive (at least⁴) two readings:

Narrow ObjO⁵: If you desire⁶ that p , and your doing A is necessary for p , then O(you do A)

Wide ObjO: O(If you desire that p , and your doing A is necessary for p , then you do A)

If Narrow ObjO is true, then from the assumption that you desire to be a successful axe-murderer, and that this requires swinging an axe through someone's body, we can conclude that you ought to swing an axe through someone's body. Surely this is not the case, so surely Narrow ObjO is false.

Wide-Scopers therefore conclude that Wide ObjO is true. This follows from the assumption that HI is true, and that Narrow ObjO and Wide ObjO are the two ways of reading HI. I'll now argue, however, that Wide ObjO has two very controversial features. One of these is an outright unintuitive consequence; the other is a feature that many philosophers would be willing to accept, but is still highly controversial. After that, we'll take a closer look at whether HI really is ambiguous in the required way, in the first place.

Symmetry

One difference between Narrow ObjO and Wide ObjO that should be immediately obvious is that Wide ObjO has a certain kind of *symmetry* that Narrow ObjO does not. Narrow ObjO says that if you desire that p , and your doing A is necessary for p , then you ought to do A . But it does not say that if you do not do A , and you desire that p , then you ought to make sure that your doing A is not necessary for p . Nor does it say that if you do not do A , and your doing A is necessary for p , then you ought to not desire that p . Wide ObjO, on the other hand, does posit a symmetry between any of these three ways of complying with its requirement. When Ronnie finds himself desiring to go dancing, and the party is the only place where there will be dancing, Ronnie can satisfy Wide ObjO by going to the party. But he can also satisfy it by ceasing to desire to dance. And he can even satisfy it—this is the crazy part—by convincing the party-throwers not to have dancing after all. For if they cancel the dancing, then going to the party won't be necessary for dancing.

Leave aside whether it is rational for Ronnie to react to his situation by ceasing to desire to dance. It *is* the right kind of thing to be a distinctively rational response to his situation, for him to go to the party. But unfortunately, convincing the party-throwers to cancel the dancing does *not* seem to be a distinctively rational response to Ronnie's situation. Whatever else we might say about it, it would be particularly odd for Ronnie, who sincerely desires to dance, to start trying to convince the party-throwers to cancel the dancing, on the grounds that he won't be able to make it. This just doesn't seem like the kind of thing that a good account of objective instrumental rationality should endorse. It is a symmetry predicted by the Wide-Scope account that is simply not sustained. Since Wide ObjO makes this prediction, it is clearly problematic.

Bizarrely, Jonathan Dancy has argued that symmetry creates a problem for the *Narrow-scope* view⁷. His argument has three parts. First, he assumes that everyone agrees that *Wide* ObjO is true. Then, he assumes that the only reason that anyone would believe *Narrow* ObjO is if it was a consequence of *Wide* ObjO. And then he uses the symmetry of *Wide* ObjO in order to derive unintuitive consequences of such a view. Dancy's argument shows that we should not think that *Narrow* ObjO is a *consequence* of *Wide* ObjO. But it does nothing to show that *Narrow* ObjO is itself problematic. For Dancy's argument *against* the *Narrow-Scope* view works by attributing to it commitment to the *Wide-Scope* view as well. Then he uses the fact that commitment to the *Wide-Scope* view engenders a symmetry, to show how this symmetry yields odd results when combined with the *Narrow-Scope* view. But obviously the symmetry only comes in when we accept *Wide* ObjO. *Narrow* ObjO by itself has no worries about symmetry.

Agent-Neutrality

Another obvious difference between *Narrow* ObjO and *Wide* ObjO is that *Wide* ObjO is committed to an eternal, agent-neutral obligation, while *Narrow* ObjO is not. To see how, let's dispense with talk about what "you" ought to do, and state *Narrow* and *Wide* in their full quantified glory⁸:

Narrow: $\forall x$ If $p(x)$, then $O_x(x$ does $A)$

Wide: $\forall x O_x(\text{If } p(x), \text{ then } x \text{ does } A)$

Wide tells us that there is something that everyone ought to do. *Narrow* tells us no such thing. As far as *Narrow* is concerned, there are only things that particular people ought to do—those people who satisfy the relevant conditions.

This difference between *Wide* and *Narrow* lies at the heart of a great controversy about the priority of agent-neutral and agent-relative obligations. Some hold that every time some individual ought to do something, it must be because there is something that everyone ought to do. These are the *Neutral-Prioritists*. But others reject this. They hold that when there is something that everyone ought to do, that is simply because *each* individual ought to do it. Indeed, the standard definition of agent-neutrality works in this way. It says that there is an agent-neutral reason to do something, just in case there is a reason for everyone to do it. Those who take this line are the *Relative-Prioritists*.

The divide between *Neutral-Prioritists* and *Relative-Prioritists* is old and deep. *Neutral-Prioritists* are happy to accept principles like *Wide*. For they think that everything that someone ought to do has to be explained by something that everyone ought to do. For them, the question of how to account for objective instrumental rationality is precisely the question of how to use an agent-neutral obligation, in order to explain how desires or ends can affect what some particular individual agent-relatively ought to do. But *Relative-Prioritists* are

not happy to accept principles like Wide. They hold that not all agent-relative “oughts” can be explained by agent-neutral ones. And they hold that agent-neutral “oughts” carry an explanatory burden. For something must explain why it is that each and every possible agent happens to bear the “ought” relation to this one particular thing. So Relative-Prioritists find principles like Wide particularly suspicious⁹.

One apparently very common kind of Relative-Prioritist view is clearly committed to Wide ObjO being false. This is the view variously known as the “Humean” Theory of Reasons, or the “Desire-Dependence” view. According to this view, all *oughts* or reasons are just like the ones accounted for by the account of objective instrumental rationality. Whenever there is a reason for someone to do something, on this view, it is because doing so promotes one of her desires. This view is not committed to holding that there are no agent-neutral obligations—it is simply committed to holding that if there are, it is *because* they are obligations that happen to be obligations for each agent, rather than vice-versa. This makes it a relative-prioritist view. According to E.J. Bond, this view was in fact “the favoured view among professional philosophers” as recently as 1983¹⁰; T.M. Scanlon recently writes¹¹ that desires are still “commonly understood” to be the sole source of reasons in this way.

There are actually *two* reasons why Wide ObjO is inconsistent with the “Humean” Theory of Reasons, but one arises simply from the fact that Wide ObjO is committed to explaining Ronnie’s case by means of a further, agent-neutral requirement. According to the “Humean” Theory, all obligations or reasons get explained in the same way as Ronnie’s reason to go to the party gets explained—by desires. But according to the Wide-Scope account, Ronnie’s reason to go to the party needs to be explained by the existence of a further agent-neutral requirement. This further requirement therefore can’t be explained in the same way as Ronnie’s reason to go to the party, because then it would have to be used to explain itself, and that would be circular. So it *can’t* be explained by a desire. Anyone who accepts Wide ObjO as an account of objective instrumental rationality, therefore, thereby rejects the “Humean” Theory of Reasons.

It has been argued on these grounds that the “Humean” Theory of Reasons is incoherent¹². But that’s silly. Such an argument employs a controversial premise—the Wide-Scope account of objective instrumental rationality. It *may* be that the Wide-Scope account is well-motivated. But at the worst, that would pose a dilemma for the “Humean” Theory of Reasons—not demonstrate it to be literally incoherent. Now it may be that the “Humean” Theory of Reasons is false. It may even be, although having thought about the matter a great deal and being consequently sympathetic to the “Humean” theory¹³ I would be quite surprised, that it is *obviously* false. But if the “Humean” theory is really “the favoured view” or even “commonly” accepted, the Wide-Scope account of objective instrumental rationality simply can’t be “uncontroversial”.

A Tangent—The Ambiguity in HI

I now turn to whether Wide ObjO is, in fact, well-motivated. Wide-Scopers typically motivate Wide ObjO by an argument from elimination. Narrow ObjO yields absurd results, so it can't be true. Therefore Wide ObjO is. But curiously, only two things made it into this argument by elimination. The justification for this is that we already know *something* about how to account for objective instrumental rationality—it is by HI. And the Wide-Scoper claims that fortunately, HI is ambiguous between the Wide and Narrow readings. Unfortunately, however, if we take this ambiguity claim seriously, it is rather implausible. For it relies on a very problematic semantics for “ought”.

The assumption that we need, in order to get the Wide-Scoper's argument going, is that “ought” expresses a relation that takes propositions for one of its objects. The ambiguity proposed by the Wide-Scoper is that in HI, the “ought” can be read as taking scope over the whole conditional, or merely over the consequent of the conditional. These are the two sentential clauses in which it figures, so if we think that “ought” takes propositions and works like a sentential operator, then this ambiguity makes sense.

On the face of it, however, “ought” does not take propositions. It takes *actions*, in some very broad sense—things that people can *do*. This is why “there is something that you ought to do” follows from “you ought to go to the store” and “going to the store is something that you ought to do” rearranges pleonastically with “you ought to go to the store”. Propositions are not things that you can *do*. So if “ought” takes propositions, then “there is something that you ought to do” should not follow from “you ought to go to the store”. Something like “there is something that you ought to make true” should follow instead. Likewise, *going to the store* is not a proposition. It is an action-type. So if “ought” takes propositions, then it is hard to see why “going to the store is something that you ought to do” should pleonastically rearrange with “you ought to go to the store”. But if “ought” takes action-types, on the other hand, then this is easy to see. On this view, these two sentences are related to one another in the same way as “Mary is left of John” is related to “John is someone Mary is left of”.

The thesis that “ought” takes propositions, as John Broome notes¹⁴, is not without linguistic evidence. The evidence for this view comes from an attractive proposal for how to understand infinitive clauses like “to go to the store”. Compare “he wants to see the Pacific” to “he wants her to see the Pacific”. On a natural view, these two sentences should be accounted for along similar lines. So on a natural view, there must be a hidden pronoun in “he wants to see the Pacific.” It must really be a little bit like, “he wants himself to see the Pacific”. Granting the existence of such hidden pronouns, it looks like the infinitive clause gives us something very proposition-like. On the other hand, for the reasons just cited, and a few others, we shouldn't get over-excited by this kind

of evidence. Compare these sentences to one like “it is wrong to murder children”. What is the hidden pronoun in “it is wrong to murder children”? For whom does it say that murdering children is wrong? Perhaps such sentences are best treated as involving some kind of generic or universal quantifier, but this is surely highly controversial, at best¹⁵. This sentence seems to pleonastically rearrange with “murdering children is wrong” and to predicate wrongness of an *action*—murdering children. Whatever kind of things we *ought* to do, plausibly they are the same kind of thing as whatever kind of things are *wrong*.

Worst of all, the thesis that “ought” takes propositions yields some intolerable predictions. It predicts that it should be at least conceptually possible that you ought that *I* go to the store. After all there is a proposition that *I* go to the store. And you are an agent. And those are the kinds of thing that the “ought” relation holds between. So all it takes for it to be the case that you ought that *I* go to the store is that you stand in the *ought* relation to this proposition. That is, that $O_{\text{you}}(\text{I go to the store})$. Now, it is certainly possible that you ought to make sure that *I* go to the store. And it is certainly possible that you ought to tell me to go to the store, and that you ought to *help* me go to the store. These are the possibilities that $O_{\text{you}}(\text{you make sure that I go to the store})$ and $O_{\text{you}}(\text{you tell me to go to the store})$ and that $O_{\text{you}}(\text{you help me go to the store})$. These possibilities all make sense. But if Broome’s view about “ought” is right, then there should be another possibility, distinct from all of these: the possibility that $O_{\text{you}}(\text{I go to the store})$. Frankly, I can’t see what this could even *be*. It sounds like a category mistake.

Broome is happy to bite this bullet. In print and in personal conversation, he has expressed his regrets that English grammar does not let us talk about such interesting possibilities as your *ought*-ing that *I* go to the store¹⁶. But that *is* simply bullet-biting. There is no such possibility for us to talk about. The claim that you ought that *I* go to the store isn’t simply an ungrammaticality. After all, the claim that you ought that *you* go to the store is also ungrammatical, but we can at least understand what it means. The claim that you ought that *I* go to the store is worse than ungrammatical, for so long as we distinguish it from each of the other things we distinguished it from above, none of us have any idea what it means. It patently manifests a category mistake. So “ought” simply can’t properly express a relation between agents and propositions. If it did, there really would be such possibilities to talk about.

I’ve been following Broome in taking the reasonable view that “ought” takes agents for one of its relata. But there is another view on which “ought” takes propositions instead of actions, but on which it is not a relation between an *agent* and a proposition. It treats “ought” like the English expression, “it ought to be the case that”, as expressing a monadic property of propositions. Those who hold this view have a different but related problem to deal with. They must give us an analysis of “*you* ought to go to the store”. On the natural version of this view, for it to be that you ought to go to the store is just for it to be the case that $O(\text{you go to the store})$. But there is an important difference

between “you ought to go to the store” and “the deficit ought to shrink” that this analysis seems not to capture. For you, unlike the deficit, are an *agent*. And most philosophers think that there is a sense in which it can be that an agent ought to do something that cannot apply to non-agents. But the deficit figures in the subject-place of O(the deficit shrinks). So it looks like if “you ought to go to the store” follows from O(you go to the store), then “the deficit ought to shrink” must follow in precisely the same sense from O(the deficit shrinks). So this view, like Broome’s, has obvious troubles making the right predictions about the sense in which an agent ought to do something. Broome’s view predicts too many things for it to be that an agent ought to do. This other view predicts too many things to qualify as “ought”-ing to do something¹⁷.

Reasons

If “ought” does not take propositions, then HI is not, after all, ambiguous between Narrow and Wide readings. HI is a conditional, conditionals are not actions, and the Wide scope reading requires that the “ought” take the entire conditional for its scope. So Wide ObjO is not an admissible disambiguation of HI. This means that if Narrow ObjO is false, then HI is false. And if HI is false, then the right way to give an account of objective instrumental rationality must be to discover what is true *instead* of HI—not to find a reading of HI on which it is true.

Now, this is something that we can understand Wide-Scopers as trying to do. On this revised reading, Wide-Scopers are not offering Wide ObjO as a *disambiguation* of HI. They are offering it as a *replacement* for HI—that is, for Narrow ObjO. It is weaker than Narrow ObjO in one relevant respect. From Wide ObjO and the assumption that you desire that *p* and your doing *A* is necessary for *p*, nothing follows about whether you ought to do *A*. That such conclusions *did* follow from Narrow ObjO was precisely what was wrong with it.

But now that we are engaged in this project, it is easy to see that there are other ways of weakening Wide ObjO in order to get this result. For example, we can replace talk about what you *ought* to do with talk about what there is a *reason* for you to do. We can assume that if you ought to do something, then there must be *some* reason for you to do it. But you have reasons to do many things that you ought not to do—even things that you *patently* ought not to do. For not all of your reasons are very good. The reasons for you to perform some particular action are a little bit like the items which appear in the “pros” column when God sits down and lists all of the pros and cons of your performing that particular action, with a view to advising you about whether to do it. Even if God always advises you one way or the other, he almost always has at something to mark in each column. So there is almost always at least *some* reason in favor of any course of action, even ones you patently ought not to take.

Now, it is reasonably obvious that even when you desire to become a successful axe-murderer, you still ought not to sharpen your axe or stake out

victims, let alone swing your axe at people. But it is less obvious that you have *no* reason whatsoever to do so. After all, the reason might simply not be very good—and we can agree that the reasons for you *not* to do these things are about as excellent as reasons come. So even if you do have some reason to do these things, we need have no worries about whether it will turn out that you ever *ought* to do them. This gives us a quite different, *Narrow-scope* way of weakening Narrow ObjO in order to avoid its unintuitive results:

Narrow ObjR: If you desire that p , and your doing A promotes p , then there is a reason for you to do A .

Narrow ObjR is clearly a Narrow-Scope account of objective instrumental rationality. But it is not at all obvious that its consequences are intolerable¹⁸. Notice that in addition to weakening Narrow ObjO by changing “ought” to “reason”, I’ve strengthened it, by replacing “is necessary for p ” with “promotes p ”. I’ll return to discuss this kind of change when we get to the discussion of theoretical rationality, in section 5.

2. Subjective Instrumental Rationality

As it turns out, different issues arise, and with more or less force, when we consider the Wide-Scoping program in the different domains in which it is applied. So in the next few sections I’m going to go through and consider each of four more such applications, in order to bring out a few more complications. The case of Ronnie and Bradley clues us in to the fact that we need an account of objective instrumental rationality. Ronnie differs from Bradley because he is keen on dancing, but Bradley is not. Now Freddie, like Ronnie, is keen on dancing. But Freddie knows something that Ronnie does not. He knows that there will be dancing at the party. Just as Ronnie differs from Bradley with respect to whether each ought to go to the party, so also Ronnie differs from Freddie. We’d be surprised if Ronnie went to the party, but not surprised if Freddie went. We’d think Freddie irrational for not going, but not so Ronnie. The difference between Ronnie and Freddie is accounted for by an account of *subjective* instrumental rationality.

Philosophers discussing instrumental reason are often not very careful to explain whether they are talking about the difference between Ronnie and Bradley, or the difference between Ronnie and Freddie. But these *are* two distinct differences, and this is important. A Wide-Scoper would have it that we can all agree that **(HI+)** you ought, if you desire that p and believe that your doing A is necessary for p , to do A . But of course this gets two readings:

Narrow SubjO: If you desire that p , and you believe that your doing A is necessary for p , then O(you do A)

Wide SubjO: O(If you desire that p , and you believe that your doing A is necessary for p , then you do A)

Narrow SubjO is twice as unintuitive as Narrow ObjO. For now there are two ways to derive crazy results from Narrow SubjO. We can assume that you have crazy desires, or we can assume that you have crazy beliefs about how to accomplish your desires. For example, you might desire to succeed in your career, and falsely believe that murdering me in cold blood and spreading my remains around your boss's office is the way to do so. But it hardly seems to follow from this that you ought to murder me in cold blood and spread my remains around your boss's office. So Wide SubjO is to be preferred to Narrow SubjO.

Symmetry

Like all Wide-Scope principles, Wide SubjO posits a symmetry between different ways in which it might be fulfilled. Freddie can satisfy the requirement posed by Wide SubjO in any of three ways. He can go to the party, or he can stop desiring to dance, or he can change his mind about whether there will be dancing at the party. But this is peculiar. Surely, concluding that there will not be dancing at the party after all is not a distinctively instrumentally rational way of responding to Freddie's situation.

We need to be somewhat careful, here. For the Wide-Scoper can say that this is indeed an irrational way for Freddie to respond—because it is ruled out by his account of *epistemic* rationality. So we don't get a problem for the Wide-Scoper merely by noticing that so far as it says, it may be rational for Freddie to change his belief. The problem for the Wide-Scoper is that if Freddie *does* respond to his situation in this way, she has to allow that though Freddie is being *epistemically* irrational, he is in fact behaving impeccably, when it comes to subjective instrumental rationality. And that is a bizarre thing to say. Surely a good account of subjective instrumental rationality should not tell us that so far as instrumental rationality goes, this kind of behavior is okay.

Freddie's case illustrates an important point. Wide-scope principles are good at predicting what is wrong with an agent *at a time*. But they are not good at predicting the rational ways for an agent to *change* her situation. Now, not all domains in which the Wide-Scoping program is applied are domains in which we are particularly interested in how it is rational for an agent to respond to her situation. For example, this is explicitly one difference between the domain of objective instrumental rationality and the domain of subjective instrumental rationality. It is *irrational* for Freddie not to go to the party, but not irrational for Ronnie not to go. Since Ronnie doesn't know anything about the party, going there is no more rational than not—even though, in some sense, it is what he ought to do, given his ends.

In Freddie's case, if he is not going to the party, something is going badly. He wants to dance, he believes that he can only dance by going to the party, and he doesn't go. If he then changes his mind about whether there will be dancing at the party, then he puts himself in a better position. He no longer has this kind of inconsistency between his aims, beliefs, and actions. So he takes himself from a worse position to a better. But despite the fact that this kind of move makes him more rational at a time, it is not a rationally permissible move. If we want an account of subjective instrumental rationality to specifically tell us something about what moves it is rational for Freddie to make, then we should be particularly sensitive to the fact that the Wide-Scope account predicts only symmetries.

Agent-Neutrality

According to Wide SubjO, there is something that everyone ought to do—to not have desires, beliefs and actions in conflict with one another in the way that Freddie's are, when he doesn't go to the party. But it's hard to see where the obligation to do this comes from. It doesn't arise because of desires, nor because of beliefs. A Narrow-Scope view can say where the obligations or reasons that it posits come from—they arise as a result of beliefs or desires. The Wide-Scope view, on the other hand, needs to posit *unexplained* obligations or reasons.

In fact, if we accept Wide ObjO and Wide SubjO, then we have to posit two distinct eternal, agent-neutral requirements in order to explain what is (distinctively) wrong with Freddie when he doesn't go to the party and what is wrong with Ronnie when he doesn't go. But on a natural view, there should be some common explanation of what goes wrong with Ronnie and what goes wrong with Freddie. They should be related in some intimate way. I'll illustrate how to give such an explanation in the next subsection.

Subjective Reasons

As in the theory of objective instrumental rationality, I'm going to suggest that Wide-Scopers weaken Narrow SubjO in the wrong way. Or at least, they don't weaken it in the only plausible way. But in this case, it is not sufficient to say that there is a *reason* for Freddie to go to the party. For though Freddie believes there to be dancing at the party, perhaps it turns out that Freddie is wrong. If there is no dancing at the party, then it follows from the account of objective instrumental rationality that there is a reason for Freddie *not* to go to the party. But it sounds more than odd to say that at least there is this much to be said for his going there anyway: at least he *believes* that there will be dancing there. That shouldn't make it into God's list of pros and cons. So the theory of subjective instrumental rationality raises a new puzzle not raised by the theory of objective instrumental rationality. If we are to weaken Narrow SubjO, we need to weaken it *more* than by changing "ought" to "reason".

I think that this shouldn't be surprising. For I don't think that we should want a distinct account of subjective instrumental rationality in the first place. Compare Ronnie and Freddie to Ryan and Bryan. Katie needs help, and that's a reason to help her. It's a reason for Ryan to help her, and a reason for Bryan to help her. But only Bryan knows that Katie needs help. Ryan is blissfully unaware. So Ryan and Bryan differ in what we can expect of them, in both the predictive and the normative senses. We can expect Bryan to help Katie, but we can't expect any such thing of Ryan. We can blame Bryan for not helping her, but we can't blame Ryan¹⁹.

It looks like Bryan differs from Ryan in exactly the same way as Freddie differs from Ronnie. There is some reason for each to do something, but only Bryan and Freddie are aware of these reasons. There is a special kind of status that you have, when you believe something that, if it is true, is a reason for you to do something. In ordinary English, in fact, we can even use the word "reason" to describe your situation. Consider the familiar case of Bernie: Bernie is at a cocktail party, holding a glass of gasoline that he believes to be a gin and tonic. Intuitively, the fact that his glass is full of gasoline is a reason for him not to take a sip. But in another sense, this isn't one of his reasons—at least, it isn't a reason that he *has*, since he is unaware of it. In this second, perfectly legitimate sense, Bernie *does* have a reason to take a sip—for he reasonably believes that his glass contains the gin and tonic for which he asked the hostess. Carefully spelled out, Bernie's case gives us cause to distinguish these two senses of the word "reason": call them *objective* and *subjective*.

On a natural view, subjective reasons are simply things that you believe such that, if they are true, they are reasons for you to do something. That seems to be what is going on with Bernie, it seems to be what is going on with Bryan, and it seems to be what is going on with Freddie. If this is right, then it follows from our account of *objective* instrumental rationality that Freddie has a *subjective* reason to go to the party that Ronnie doesn't have. For Freddie, but not Ronnie, believes that there will be dancing at the party, and this is the *objective* reason for both of them to go there.

Taking this very natural view of the matter commits us to an even weaker account of subjective instrumental rationality:

Narrow SubjSR: If you desire that p , and you believe that your doing A is necessary for p , then you have a subjective reason to do A .

But on this view, Narrow SubjSR is not a distinct theoretical posit, needed in order to explain what is going on in the case of subjective instrumental rationality. It falls neatly out of our already-existing Narrow-Scope account of objective instrumental rationality, and our very natural account of the relationship between the objective and subjective senses of the word "reason", which we are independently forced to acknowledge.

3. The Role of Conscience

Wide-Scoping is also commonly employed to explain the relationship between your conscience and what you ought to do. According to the Wide-Scoper, we can all agree that in some sense or other (**Consc**) you ought, if you believe that you ought to do something, to do it. You should, that is, let your conscience be your guide. But as always, the Wide-Scoper holds that this claim is ambiguous:

Narrow ConscO: If you believe that O(you do *A*), then O(you do *A*)

Wide ConscO: O(If you believe that O(you do *A*), then you do *A*)

Now, Narrow ConscO yields some unfortunate results. From it, it follows that you are infallible with respect to what you ought to do. And surely that is false. Surely you can be mistaken about what you ought to do. So Wide ConscO is definitely to be preferred to Narrow ConscO.

Symmetry

But let us not get ahead of ourselves. Wide ConscO, too, has its problems. For one, there are *two* ways to comply with Wide ConscO. You can comply with it either by doing what you believe you ought to do, or by changing your mind about what you ought to do. But surely there is a relevant asymmetry, here. After all, we have a special name for the distinctive vice of changing your mind about what you ought to do, simply so that you don't have to do it. It is called *rationalization*. The whole point of conscience being your guide is that changing your beliefs about what you ought to do simply in order to avoid doing it is *not* an acceptable way to proceed.

Agent-Neutrality

Like the other Wide-Scope accounts, Wide ConscO works by positing a basic, eternal, agent-neutral requirement rationally binding on every agent, no matter what they are like. As with all of the others, it offers no explanation of this requirement. Narrow-Scope accounts can explain the obligations or reasons that they postulate. After all, these obligations or reasons only exist given a certain condition—so we can use that condition to explain them²⁰. But not so for the Wide-Scopers.

Subjective Reasons Again

The agent-neutral requirements postulated by the Wide-Scoper involve even more commitment, once we see that the requirements postulated by Wide ObjO,

Wide SubjO, and Wide ConscO are all distinct. I propose, however, that we can reject Consc altogether, at least if we understand both “oughts” in the same *objective* sense. And as in the other cases, I suggest that our theory can be better—and more economical—if we find a different way to weaken Consc. As with the theory of subjective instrumental rationality, I suggest that the necessary weakening involves the notion of a *subjective* reason:

Narrow ConscSR: If you believe that O(you do *A*), then you have a subjective reason to do *A*.

On the natural view that I will suggest, Narrow ConscSR simply falls out of the account of what subjective reasons are, from section 2. This is because on a natural view, the fact that you ought to do *A* is a reason for you to do *A*.

A number of philosophers have recently rejected this natural view. They claim that the fact that you ought to do *A* only *reports* the existence of *other* reasons for you to do *A*—it is not itself a reason for you to do *A*. Their argument is twofold: first, it can’t be the case that you ought to do *A* unless there is some *other* reason for you to do *A*. And second, the fact that you ought to do *A* shouldn’t be weighed *separately* from these other reasons, in determining whether you ought to do *A*. It doesn’t carry any extra weight. I agree with both of these claims. But weighing reasons is complicated. Perhaps two things can both be reasons, even though they shouldn’t be weighed separately. For example, in ordinary English, we can say that the fact that there will be dancing at the party is a reason for Ronnie to go there. But in ordinary English we can also say that the fact that Ronnie is keen on dancing is a reason for Ronnie to go to the party. Surely, if these are both reasons for Ronnie to go to the party, they shouldn’t be weighed separately. Now some hold that this is an argument that ordinary English speaks falsely on these counts, and these are really the same reason. But I hold that we could just as well say that adding up the weights of reasons is more complicated than simply placing weights on a scale. If that is right, then we can say that the fact that you ought to do *A* itself counts as a reason for you to do *A*.

Indeed, given our account of the relationship between objective and subjective reasons, this is precisely what we should say. For consider the case of John²¹. John loves successful surprise parties thrown in his honor, but hates all other parties—most of all unsuccessful surprise parties thrown in his honor. In the next room, all of John’s friends are waiting, ready to surprise John with a party of which he so far has no clue. There is an excellent reason for John to go into the next room—that a successful surprise party is waiting for him. But you could never give him this reason, because that would make the reason itself disappear. So instead you merely tell him that he *ought* to go into the next room. Does John have a subjective reason to go into the next room? Well, we can expect him to go, and it would be irrational of him not to go there. If he goes, we won’t say that he went for no reason at all—we’ll say that he went because he

ought to. If this is right, then John has a subjective reason to go into the next room by believing that he ought to go. And so by our account of subjective reasons, the fact that he ought to go into the next room must itself count as an objective reason for him to do so. And granting this, Narrow ConscSR falls out immediately from our account of subjective reasons.

4. Promises

A very old application of Wide-Scoping is to promises. This domain sheds some light on another possible motivation for Wide-Scoping, as well as on the kind of commitment I've been discussing under the heading of "agent-neutrality". Intuitively, the issue is this: Al promises Rose to meet her for lunch. Fortunately for Al, if something comes up, Rose has the power to excuse him from this promise. But if she doesn't, then something is amiss if Al doesn't show up for lunch. According to the Wide-Scoper, we can all agree that in some sense or other (**Promise**) you ought, if you promise *Y* to do *A* and *Y* doesn't excuse you, to do *A*. And as always, the Wide-Scoper claims that we can read this in more than one way:

Narrow PromiseO: If you have promised *Y* to do *A*, and *Y* has not excused you from doing *A*, then O(you do *A*)

Wide PromiseO: O(If you have promised *Y* to do *A*, and *Y* has not excused you from doing *A*, then you do *A*)

The normal Wide-Scoping puzzle would be that Narrow PromiseO apparently can lead to some funny consequences. What if you promise someone to commit some foul deed? Ought you to commit the foul deed? Or what if you make conflicting promises? Ought you to do both? To avoid these results, the Wide-Scoper would have us to prefer Wide PromiseO to Narrow PromiseO.

But in fact, this is not usually what motivates Wide-Scoping about promising. What usually motivates Wide-Scoping about promising is a much simpler observation: that there is only one situation in which something is really going wrong with Al. It is the situation in which he makes his promise, Rose does not excuse him from it, and he fails to show up for lunch. The possibilities are depicted on the following table:

	<i>Al shows up for lunch.</i>	<i>Al doesn't show up.</i>
<i>Doesn't give permission.</i>	Okay	! something amiss !
<i>Rose gives permission.</i>	Okay	Okay

The Wide-Scoper's natural idea about Al and Rose's case is that what we need is some way of *ruling out* the situations in which something goes amiss. So that is

precisely what she does. She postulates a special requirement that says, “don’t be like that,” demonstrating the cases in which something goes amiss because an agent makes a promise, is not excused, and doesn’t keep it.

This kind of motivation can be supplied for Wide-Scoping in other domains. In the theory of objective instrumental rationality, we want to know what is amiss with Ronnie, when he desires to dance, going to the party is necessary for dancing, and he doesn’t go to the party. So the Wide-Scoper postulates a special requirement that rules out precisely those kinds of situation. In the theory of subjective instrumental rationality, we want to know what is amiss with Freddie, when he desires to dance, believes that there will be dancing at the party, and doesn’t go to the party. So the Wide-Scoper postulates a special requirement that rules out precisely those kinds of situation. In theorizing about the role of conscience, we want to know what is amiss with you, when you believe that you ought to do *A*, but don’t do *A*. So the Wide-Scoper postulates a special requirement that rules out precisely those kinds of situation. In this way, Wide-Scoping gains incredibly elegant solutions to each of these problems—always by postulation of some new eternal, agent neutral requirement, irreducible to any of the others. Elegant solutions, at the cost of unexplained, basic, agent-neutral requirements. This is a distinct kind of motivation for Wide-Scoping.

Symmetry

The Wide-Scope account of promising postulates a single requirement ruling out precisely those situations in which something is amiss. So to satisfy this requirement, all that Al has to do is to get out of this situation. Given that he’s already promised to meet Rose for lunch, he can’t change that fact. But he *can* solicit her permission not to show up. Rose has the power to *excuse* Al from showing up for lunch. If Al convinces her to do so, on the Wide-Scope view, he is *satisfying* all of his relevant obligations.

But that is clearly wrong. There is at least one obligation that Al is precisely *not* satisfying by soliciting Rose’s permission not to show up for lunch. It is an obligation that he is getting *out* of. And this is the obligation to show up for lunch. Though Rose has the authority to *dismiss* Al’s obligation to show up for lunch, that doesn’t change the fact that he *does* have such an obligation. It is an obligation that he *satisfies* by showing up for lunch, but merely *escapes* by getting Rose’s permission not to. Either way, he is only *violating* the obligation if he both does not show up for lunch and does not get her permission.

The problem is that the table only shows us when Al has fallen astray of some obligation *or other*. It doesn’t tell us anything about which, if any, obligations he is *satisfying*, or which, if any, obligations he is *escaping*. Unlike the asymmetries that we’ve diagnosed so far, this is not an asymmetry that only exists when we look at how we want or expect an agent to behave *over time*. In Al’s case, asking Rose’s permission not to show up for lunch is in perfect

keeping with the rules governing promises. It's a *perfectly* acceptable thing to do. But after he has done it, the correct way to describe his situation is that his obligation has *gone away*—not that he has satisfied it.

Agent-Neutrality

The Wide-Scope account of promising therefore deals with Al's case a little bit *too* neatly. The mere fact that something is going amiss in a large class of situations doesn't demand that there be any requirement in particular that rules all of those situations out. In Al's case, Narrow PromiseO succeeds in ruling out all and only the same situations as does Wide PromiseO. But Narrow PromiseO rules out these situations on a case-by-case basis. When Al doesn't get permission from Rose and fails to show up for lunch, the obligation that he is violating is his obligation to meet Rose for lunch. And we know where that came from—he promised her to. Likewise, when Sylvia fails to repay her loan and is unexcused, the obligation she is violating is her obligation to repay her loan. And Narrow Promise tells us where this obligation came from, as well—it arose as a result of the promise that she made when she asked her parents for money.

The promising case is therefore good for illustrating how a Narrow-Scope account can rule out all of the same situations as a Wide-Scope account, but without positing symmetry, and without committing to a single and unexplained requirement that rules those situations out specifically. And this further illustrates the nature of the commitment of Wide-Scopers that I've been discussing under the heading of "agent-neutrality".

Obligations

In the case of promising, I don't think that much needs to be done in order to amend Narrow PromiseO. This is because the consequences of Narrow PromiseO, interpreted correctly, are not particularly unintuitive. In fact, in this case we know from the asymmetry of satisfying obligations and escaping them that we *need* at least one Narrow-scope principle in order to distinguish the case in which Al gets excused from that in which he keeps his promise. And once we have that principle, any Wide-Scope account would be superfluous, and can be done without.

Still, philosophers are often picky about the word "ought". It is often held that it can't be the case that you ought to do *A* and that you ought to do *B*, if doing *A* and doing *B* conflict. This view follows, in fact, from the principle that *oughts* aggregate across conjunction, and the common view that "ought" implies "can". But it is clearly possible to make conflicting promises. I'm not sure what to think of either of these two principles, but if this is really how "ought" works, then fortunately there is a very simple way of weakening Narrow PromiseO to avoid this result:

Narrow PromiseOb: If you have promised *Y* to do *A*, and *Y* has not excused you from doing *A*, then you are under an obligation to do *A*.

Narrow PromiseOb posits an obligation for you to do *A*, when you promise to do *A*. On this view, you may have conflicting obligations, but it does not follow from the fact that you are under an obligation to do *A*, that you ought to do *A*. This only follows if you are under no conflicting obligations. So it still turns out that it can't be the case that you ought to do *A* and that you ought to do *B*, if doing *A* and doing *B* conflict.

5. Epistemic Rationality

Philosophers defending Wide-Scope accounts of one or another of the domains discussed above often appeal to the domain of epistemic rationality as a case in which Wide-Scoping obviously applies, in order to gain credibility for their views. Intuitively, the case is this: Phil believes that *p*, and that if *p* then *q*. But he also disbelieves that *q*. Patently something is amiss with him. At least four different kinds of account might be offered:

- Narrow Bf(wk)O:** If you believe that *p* and that if *p* then *q*, then O(you don't disbelieve that *q*)
- Wide Bf(wk)O:** O(if you believe that *p* and you believe that if *p* then *q*, then you don't disbelieve that *q*)
- Narrow Bf(sg)O:** If you believe that *p* and that if *p* then *q*, then O(you believe that *q*)
- Wide Bf(sg)O:** O(if you believe that *p* and you believe that if *p* then *q*, then you believe that *q*)

The weak (wk) accounts govern only whether you should disbelieve that *q*. The strong (sg) accounts govern whether you should actually *believe* that *q*. Here I assume that one can withhold belief from a proposition—take no opinion about it. This is neither believing nor disbelieving.

Wide Bf(wk)O successfully rules out the worst set of cases: those in which you actually have contradictory beliefs. Something is amiss with you if you have contradictory beliefs, whether or not you realize it. But Wide Bf(wk)O doesn't tell us very much about how it is rational for you to react to your situation. For it doesn't even rule out the situation in which Phil believes that *p*, believes that if *p* then *q*, is actively wondering whether *q*, and still has no opinion about whether *q*. And that, surely, is a situation in which it is not epistemically rational for Phil to find himself.

Yet Wide Bf(sg)O seems to rule out too much. For it rules out situations in which Phil believes that *p* and believes that if *p* then *q* but has never put these two thoughts together, which is what explains why he has never formed an

opinion about whether q . And that is clearly ruling out too much. It is not epistemically irrational at all to be like that. Nor is it rationally required for you to have a deductively closed set of beliefs.

Symmetry

Wide Bf(sg)O also predicts a tolerable, if slightly surprising, symmetry. It predicts that if you find yourself believing that p and that if p then q and having no opinion whatsoever about q , one epistemically rational way for you to respond to your situation is to cease to believe that p . This is at least initially somewhat surprising. But I don't think that it is intolerable²². A worse symmetry prediction comes to light when we try to repair Wide Bf(sg)O in order to solve the problem just posed for it in the last subsection.

The problem is that Wide Bf(sg)O tells us that *too* much is wrong with Phil—it tells us that something would be wrong with Phil even if he wasn't bearing his beliefs fully in mind or wondering whether q . So on the face of it, what we want is to somehow incorporate the facts that Phil is wondering whether q and that he is paying attention to his relevant beliefs. The Wide-Scope way to do this is simple²³:

Wide Bf(sg)O*: O(if you believe that p and you believe that if p then q and you are paying attention to these beliefs and you are wondering whether q , then you believe that q)

But there are *five* ways to comply with Wide Bf(sg)O*. For example, Phil can comply with it by ceasing to wonder whether q or by ceasing to pay attention to his relevant beliefs. Now, it may be that, considered at a time, Phil is being more rational to still wonder whether q if he is not paying explicit attention to the beliefs that entail it. So statically speaking, ceasing to pay attention to these beliefs can take Phil from a position where more is amiss with him, to a position in which less is amiss with him. It can make him better off, rationally speaking, if we are only making rational assessments of him at particular times. But it hardly follows that this is a rational thing for Phil to do! On the contrary, this looks like a paradigm of epistemic irrationality. So Wide Bf(sg)O* makes a prediction of symmetry that seems to be unfulfilled. Or, put differently, this is a respect in which the domain of theoretical rationality calls for an *asymmetry* for which the Wide-Scope account, by itself, is unable to account.

The symmetry problems for the Wide-Scope account of epistemic rationality get much worse, when we try to expand the account, in order to deal with cases of non-deductive inference. On the face of it, epistemic rationality does not only have something to do with how we deal with deductively valid arguments. It also has something to do with how we treat inductive evidence. But Wide-Scope accounts are designed merely around this special case. They become much

more problematic when we try to expand our account in order to deal with inductive evidence.

Suppose, then, that you believe that Carrie is 99% reliable, and that Carrie said that p , but are undecided whether p . If you have no other reason to believe that $\sim p$, there is something odd about you not going on to believe that p , in this case. There is no *strict* requirement here, of course, even when you are bearing both of these supporting propositions fully in mind. But most belief formation is not deductively valid. It is based on evidence that is less than fully conclusive, but that doesn't mean that there aren't any questions about which ways of proceeding are epistemically rational, and which are not.

The Wide-Scoper about epistemic rationality cannot simply say that there is a strict requirement that you not be in all three of these states at the same time. For once the evidence is less than conclusive, the requirement must also be less than strict. Although it makes sense to conclude that p on the basis of the belief that Carrie is 99% reliable and that Carrie said that p , you can in all rationality believe these things but believe that $\sim p$ —for example, if you saw that $\sim p$ with your own eyes. If the Wide-Scoping program is to have any general applicability, therefore, the “requirement” which forbids being in all three of these states at the same time must be less than strict—it must be merely one of degree.

Imagine, then, that the Wide-Scoper tells us that when a set of beliefs is *unlikely* to be true together, there is a *reason* not to believe all three of them which varies in strength according to *how likely* they are not to be all true together²⁴. This is an unpromising tactic, since it makes symmetric predictions about the rationality of giving up each belief. But in fact, “Carrie is 99% reliable” and “Carrie said that p ” are much *better* evidence that p than “ $\sim p$ ” and “Carrie said that p ” are for “Carrie isn't 99% reliable”. Posing only a slack requirement not to believe all three therefore loses track of this important element of the structure of the relationship between these three propositions.

If Wide-Scoping is not the right way to account for *inductive* applications of epistemic rationality, however, then it is highly implausible that it is the right way to account for *deductive* applications of epistemic rationality. For on the face of it the deductive case is merely a special case of a much more general question about epistemic rationality. The same point applies in the case of instrumental rationality. If Ronnie is keen on dancing, then the fact that there will be dancing at the party can give him a reason to go there, even if going there isn't *necessary* for him to dance. The Wide-Scope accounts of objective and subjective instrumental rationality, like the Wide-Scope account of epistemic rationality, are designed to fit the extreme case, in which an action really is or is believed to be *necessary* for an end. But instrumental rationality has a much broader scope. It can be affected in many ways by relationships between actions and ends that fall far short of necessity. If Wide-Scope accounts aren't well-suited to explain what goes on in the less strict cases, then they can't be the right account of the strict cases. For the strict cases are surely simply a limiting case of the less strict ones.

Agent-Neutrality

Phil believes that p , that if p then q , and is wondering whether q . A Narrow-Scoper can say what reason Phil has to believe that q : after all, he believes that p and that if p then q . This is what gives him reason to believe that q , because that p and that if p then q are the right kinds of thing, if true, to be evidence that q . But according to the Wide-Scope view, evidence and reasons for belief have nothing to do with epistemic rationality. Epistemic rationality involves complying with a single, eternal, agent-neutral requirement not to have inconsistent beliefs. Insofar as he is complying with the Wide-scope requirement, when Phil forms the belief that q in this situation, the reason for which he believes q is not that p . Nor is it that p and that if p then q . It is that epistemic rationality requires not believing that p and that if p then q and wondering whether q and bearing these beliefs in mind and not believing that q . So if the Wide-Scope account is right, then agents shouldn't believe for the reasons that are their evidence. They should believe for some non-evidential reason that is eternal and agent-neutral and unexplained.

Prima Facie Reasons

As in each other case, I think that so long as we are in the business of offering *alternatives* to the Narrow O principles, rather than in looking for disambiguations of them, there is a viable—and perhaps preferable—alternative to Wide-Scoping. Instead of weakening Narrow Bf(sg)O by widening the scope of the “ought”, we can weaken it by replacing the “ought” with a weaker normative concept.

As suggested by the last subsection, I hold that we should think of the account of epistemic rationality on a par with the account of *subjective* rationality. As in the practical case, it looks like we can draw a parallel between *objective* and *subjective* senses of the words “reason” and “evidence”. If Carrie is 99% reliable, then the fact that Carrie said that p is evidence that p . It is a reason to believe that p . But if we don't know that Carrie has said that p , then this isn't evidence that we have. It isn't among our reasons. On the other hand, if we do believe that Carrie said that p , then we do have evidence that p —we have some reason to believe that p . And we can have this reason, even if we turn out to actually be mistaken about whether Carrie said that p . Whether Carrie actually said that p makes no difference as to how rational it is for us to conclude that p on the basis of our *belief* that she did.

If this is right, then the connection between beliefs and epistemic rationality should be much like the connection between beliefs and instrumental rationality. It should be²⁵ that believing something that is evidence that p , if it is true, is how to have a *subjective* epistemic reason to believe that p . And then we should say that epistemic rationality just has to do with what you have subjective epistemic reason to believe:

Narrow BfgenSR: If you believe that p and p is, if true, evidence that q , then you have a subjective epistemic reason to believe that q .

This *general* account yields as a trivial consequence the account that we were looking for:

Narrow Bf→SR: If you believe that p and that if p then q , then you have a subjective epistemic reason to believe that q .

Unfortunately, however, in the epistemic case we clearly need to be yet a little bit more careful.

The problem is this. You may believe something that, if true, is evidence that q . But you may also believe something that would *defeat* that evidence. For example, you may believe that Carrie has just said that p . But you may also have it on good authority—from Carrie herself, say—that she was simply joking. Though by itself, the fact that Carrie just said that p is the right kind of thing to be evidence that p , you believe something that cancels the force of this reason. Intuitively, we would say that you have no reason at all to believe that p , in this situation.

So if we want to keep a Narrow-Scope account, we need to weaken our principle yet further. We should distinguish between *prima facie* and *pro tanto* reasons. You have a *pro tanto* reason to believe that q just in case you ought to believe that q , if you have no countervailing evidence. Likewise, you have a *prima facie* reason to believe that q just in case you have a *pro tanto* reason to believe that q , so long as your reason is not completely undermined. A *pro tanto* reason can be partially undermined, but if it is *completely* undermined, then we say that it is merely a *prima facie* reason.

The distinction between partial undermining and complete undermining is important. On a standard example, you see Tom Grabit come out of the library, pull a book from under his shirt, and scurry off. This is reason for you to believe that Tom just stole a book from the library. But if Tom has a twin brother Tim from whom you could not distinguish him, that can undermine your reason to believe that Tom just stole a book. For all you know, it could have been Tim. Yet your reason is not completely undermined in this case. This is easy to see, by observing that it can still get *worse*. For Tom and Tim might have a third identical sibling, Tam, whom you can visually distinguish from neither. If so, then your reason would be still worse. So it can't have been completely undermined. This contrasts with what happens if you look to the side, and notice that a movie director has his camera focused on Tom, and is just saying "Cut! One more take, Tom." If you see that, then it completely undermines your visual evidence that Tom stole a book. Such a reason would be merely *prima facie*.

The plausible Narrow-Scope account of epistemic rationality replaces Narrow BfgenSR with Narrow BfgenPF—a principle exactly like it, but with the modifier that the reasons that it invokes are merely *prima facie*.

6. The Scope of Instrumental Reason

In each domain in which Wide-Scoping is applied, the issues are somewhat different. Plausibly, in each domain the debate between Wide-Scopers and Narrow-Scopers should be adjudicated on its own grounds. And in every domain, as I've noted at a few points in footnotes, it is possible to introduce further, mixed, theories, with some Wide-scope and some Narrow-scope elements. This complicates issues in ways that I haven't had opportunity, here, to explore. Yet some issues have served as common threads to each domain of inquiry. The first is that the Narrow principles allow "detaching"—if we take as a premise the antecedent of their conditionals, we can simply apply *modus ponens* and deduce results that are largely unacceptable, at least if we restrict our attention to a relatively *strict* normative concept like that expressed by "ought". This is the original motivation for Wide-Scoping.

Another issue that raised its head in every domain was the symmetry predictions of the Wide-Scope accounts. Every Wide-scope principle yielded predictions of symmetry, and many of these failed to be substantiated. Some of these failures were worse than others. Some had a bearing on whether Wide-Scoping can plausibly be generalized to less strict cases, some a bearing on dynamic questions about rationality over time, and some a bearing on simple distinctions applicable at a time, like that between whether an obligation has been *satisfied*, or merely *escaped*.

The third issue that returned over and again was the commitment of Wide-Scope accounts to unexplained, eternal, agent-neural requirements that can be avoided by Narrow-Scope accounts. Again, the issues were different in different domains. Such requirements are *controversial*, because on one common way of thinking about the relationship between agent-relativity and agent-neutrality, agent-neutral requirements need to be explained by agent-relative ones. In the account of objective instrumental rationality, we saw that this feature made Wide-Scoping unacceptable to adherents of a supposedly widespread view about reasons—that *all* reasons must be explained by an account of objective instrumental rationality.

But the commitment to these requirements is controversial in its own right. For it involves postulating distinct requirements in order to account for each domain. But I've shown how to subsume the accounts of objective and subjective instrumental rationality under one account, and how to account for epistemic rationality and for the role of conscience merely by offering an account of the relationship between objective and subjective reasons. These accounts didn't work by postulating new and unexplained requirements—they worked simply by trying to *make sense* of a group of normative concepts weaker than that expressed by "ought". I haven't had the space, here, to flesh out the details of any of the Narrow-Scope alternatives that I've offered, but I hope that I've set out just enough that they can be seen as viable alternatives which fit into an attractive context.

Since much has recently been made of it, in particular, allow me to return to the Wide-Scope account of objective instrumental rationality. I want to explore yet one more way in which Wide-Scoping can be and ought to be *controversial*. I take it that Ronnie and Bradley differ with respect to what they have reasons to do. In particular, the fact that there will be dancing at the party tonight is a reason for Ronnie to go there, but not a reason for Bradley to go there. And this difference between Ronnie and Bradley seems to be due to a difference in what they like or desire or care about.

The “Humean” Theory of Reasons, as I indicated, holds that all reasons are explained in the same way as Ronnie’s. It is a natural view to hold, if you think that at *some* level, all reasons must be explained in the same way. For once you grant that Ronnie’s reason needs to be explained by his desire, then the challenge is on to explain just how this could work, so that a desire is required to explain Ronnie’s reason, but that is a non-essential part of the uniform explanation of reasons generally.

But as I indicated in section 1, the “Humean” Theory of Reasons runs into trouble in more than one way, when it comes to the Wide-Scope account of objective instrumental rationality. The argument from Jean Hampton discussed in section 1 is only one such way. Stephen Darwall offers a distinct argument that the “Humean” Theory of Reasons is doomed to incoherence. The argument is this: 1) the correct account of objective instrumental rationality is Wide-scope. 2) From Wide ObjO and “you desire that *p*” it does not follow that you have a reason to do what promotes *p*. 3) To derive this conclusion from Wide ObjO, you need to adduce the premise that you ought to desire that *p*²⁶. So 4) there can only be reasons for agents to do particular things, such as go to the party, if there are antecedent reasons for them to have certain ends. Therefore 5) on pain of regress, not all reasons can be explained in this way. And so the “Humean” Theory of Reasons, since it claims that they can, is incoherent.

Darwall doesn’t use the fact that Wide ObjO itself posits an unexplained agent-neutral requirement in order to create trouble for the “Humean” theory. He merely uses the fact that it doesn’t allow us to “detach” in order to create the trouble. So a view like the “Humean” Theory of Reasons *needs* to make sense of a Narrow-Scope theory of objective instrumental rationality. It needs an explanation of Ronnie’s reason that will not advert to a further reason, on pain of regress. As indicated, I hold that we can do so, if we remind ourselves of the difference between saying that there is a *reason* for someone to do something, and saying that she *ought* to do it. The former is a far weaker claim.

If Darwall and Broome and others are right, and the Wide-Scoping program truly is uncontroversial, then that does it for the “Humean” Theory of Reasons. But this argument of Darwall’s raises its own bit of trouble for the Wide-Scoping program. As I understand the “Humean” theory, it is motivated by an attempt to explain where merely agent-relative reasons come from. Wide-Scoping is part of a general picture that takes agent-neutral reasons for granted, as not requiring any particular kind of explanation. It tells us something about

what is going on in Ronnie's case. But if Darwall's argument works, then Wide-Scopers can only believe that there is a reason for Ronnie to actually go to the party, if they believe in an antecedent reason for Ronnie to desire to dance.

But now we get to the heart of the matter. Is this reason to desire to dance one that is a reason for Bradley as well? If it is, then we haven't succeeded in explaining why Ronnie has a reason to go to the party that isn't a reason for Bradley to go there. But if it is not a reason for Bradley, then the question about where merely agent-relative reasons come from simply retreats another step. According to the Wide-Scoper about objective instrumental rationality, they don't arise as a result of differences in desires. But where *do* they come from? On the face of it, wherever we think merely agent-relative reasons come from, that is going to be a domain in which it looks like how things are for some agent can have an effect on what she has reason to do. In short, it is the kind of domain in which Wide-Scopers tend to apply their theories.

If Darwall's argument really works against the "Humean" Theory of Reasons, then it is an equally good argument that Wide-Scoping has to stop *some-where*. Since there really are reasons that are merely agent-relative, they have to arise *somehow*. Some differences between agents have to actually lead to a difference in what people ought or have reasons to do—one that can eventually explain the difference between Ronnie and Bradley. The "Humean" simply figures that we might as well stop at the beginning. For the sophisticated "Humean", *narrowing* the scope of instrumental reasons is just the trick in order to plausibly *expand* the scope of instrumental *reason*. And if very much at all of this paper is correct, then it's not at all obvious that this project is crazy or "confused" or narrow-minded. Whether any Wide-Scope account is correct or not should still be very much a live issue²⁷.

Notes

1. Strictly speaking, this assumption is made only by certain interpretations of deontic logics. On some interpretations, deontic logics only study the logic of "it ought to be the case that *p*", which need not be assumed to have any connection whatsoever to the "ought" of "you ought to go to the store".
2. I don't know how early the wide-scope view was first formulated, but it is possible that Sidgwick held the view: "When (e.g.) a physician says, 'If you wish to be healthy you ought to rise early', this is not the same thing as saying 'early rising is an indispensable condition of the attainment of health.'...[I]t is not merely this relation of facts that the word "ought" imparts: it also implies *the unreasonableness of adopting an end, and refusing to adopt the means indispensable to its attainment*" (1981, p I iii 4, italics added). More recently, it has been offered as obvious or defended in, for example, Hare (1971), Hill (1973), Greenspan (1975), Darwall (1983), Gensler (1985), Hampton (1998), Broome (1999), Dancy (2000), and Wallace (2001).
3. According to Stephen Darwall, the Wide-Scope view is "uncontestable" (1983, p 15). According to John Broome, it is "an elementary and widely recognized

- point, but also one that is widely ignored” (1999, p 410). Broome calls disagreeing with Wide-Scope views a “confusion”.
4. Once we start having conditionals with conjunctive antecedents, we might think that we should be able to get readings for each combination of the conjuncts outside of the scope of the “ought”. For example, here we might expect to get “If you desire that p , then O(If your doing A can bring about p , then you do A)” and “If your doing A can bring about p , then O(If you desire that p , then you do A)”. If we really think that “ought” works in this way, then it looks like all of these should be candidates that we should have to worry about.
 5. My convention from here on will be to name principles by whether they are Wide or Narrow, an abbreviation for the domain for which they are supposed to provide an account, and a letter for the normative concept that they employ. So: “O” for *ought*, “R” for *reason*, and so on.
 6. I’m going to conduct the discussion from here on as if the question is what *desires* have to do with what you ought to do. If you think the answer is “nothing”, then please feel free to substitute “end” or “intention” wherever appropriate.
 7. See Dancy (2000, p 70–76).
 8. In stating these principles, I am throughout following Broome in assuming that “ought” expresses a *relation* between an agent and a proposition. So I assume that the “O” has always an implicit index. In order to state the quantified version of the principles, however, we need to make the index explicit, so that’s what I’ve done here. See Broome (1999, p 399).
 9. For further discussion, see my “Reasons and Agent-Neutrality”.
 10. Bond (1983, p 3).
 11. Scanlon (1998, p 37).
 12. On a very natural reading, this is the central argument of Jean Hampton in *The Authority of Reason*. Hampton, however, isn’t particularly clear on this point. I can generate two other, quite different readings, of what her central argument might be. See Hampton (1998).
 13. I defended a version of the “Humean” theory in my dissertation, *Slaves of the Passions*.
 14. As Broome notes (1999).
 15. See my paper, “Reasons and Agent-Neutrality”, for further discussion of the philosophical (not semantic) reasons why this treatment of “wrong” should be so highly controversial.
 16. Broome (1999), and also in personal conversation.
 17. Nevertheless, I am going to continue to use Broome’s notation in stating the Wide-Scope views. Many (but not all) of the Wide-Scope views can plausibly be intelligibly restated without the assumption that “ought” takes propositions. By rejecting this assumption, we simply lose the *motivation* for thinking that we are somehow *preserving* the truth of HI. And so we lose the motivation for not including other kinds of account in the argument by elimination for the Wide-Scope view.
 18. See my paper, “Weighting for a Plausible ‘Humean’ Theory of Reasons” for an extensive discussion of whether Narrow ObjR is still unacceptably strong.
 19. Here I ignore for simplicity the complication that his ignorance might turn out to be culpable.

20. This is exactly what Kant says about why hypothetical imperatives are easier to understand than categorical imperatives: “On the other hand, the question as to how the imperative of morality is possible is undoubtedly the only one requiring a solution. For it is not at all hypothetical; and hence the objective necessity which it presents cannot be based on any presupposition, as was the case with the hypothetical imperatives” Kant (1997, 4:419). It is strange, then, that many Wide-Scopers interpret Kant as sharing their view about objective instrumental rationality. I argue that Kant is not a Wide-Scoper in my (forthcoming).
21. Thanks to Nate Williams for this case, even though he wouldn’t approve of this use of it.
22. It *may* be, for example, that the reason that you withhold belief from q is not that you have no evidence either way, but rather that you have *too much* evidence *both* ways, and so you can’t decide the issue. If you have enough evidence both ways, you might think that no evidence that p could possibly be robust enough to decide the issue whether q . And so the evidence that $\sim q$ might therefore be sufficient to override your evidence for believing that p even without being sufficient to justify your believing $\sim q$. This seems like a perfectly coherent scenario to me, so perhaps this symmetry prediction is quite alright, initially surprising as it might be.
23. There is a natural way of trying to give a *mixed* wide/narrow account in order to fix this problem:

Mixed Bf(sq)O: If you are paying attention to your belief that p and your belief that if p then q , and you are wondering whether q , then O(if you believe that p and you believe that if p then q , then you believe that q).

The mixed account is Narrow-Scope with respect to the paying-attention and wondering, but Wide-Scope with respect to the believing. But unfortunately, it is hard to make sense of the mixed view in a way that respects the fact that it aspires to be Wide-Scope with respect to the believing. For paying attention to a belief isn’t independent from having that belief—it seems to *require* having that belief. So once we move these conditions outside of the scope of the “ought”, we’ve already moved the believing condition outside of the scope of the “ought”, and then we don’t have a mixed view at all, but only a Narrow-Scope view.

24. This, I assume, is the obvious way for the Wide-Scoper to tackle this issue. I don’t assume that it’s the *only* way for such a view to tackle the issue, for it’s not my purpose to refute the Wide-Scope view. I’m only trying to demonstrate that such a view has some hard questions to answer, before we can agree that it is an *adequate* account of the domain, let alone that it is uncontroversial.
25. More or less. Obviously we need complications to explain how background knowledge can affect what counts as evidence for you, and so on.
26. This assumption is controversial, and may be false. Given a suitable version of the principle that “ought” implies “can” and the assumption that it is sometimes practicably impossible to give up a desire, it may turn out that the only way to comply with Wide ObjO is to take the means to such a desire. Or it may be that there is some special requirement not to give up ends, once you have them. These

would be ways in which 3) would be false, but they would also be ways in which 2) would be false. See Greenspan (1975) for an excellent discussion of the first kind of possibility.

27. Special thanks to Jim Pryor, Ant Eagle, Gideon Rosen, David Sussman, John Hawthorne, Stephen Darwall, John Broome, and Stephan Leuenberger.

References

- Bond, E.J. (1983). *Reason and Value*. Cambridge: Cambridge University Press.
- Broome, John (1999). "Normative Requirements." *Ratio*. 12(4): 398–419.
- Dancy, Jonathan (2000). *Practical Reality*. Oxford: Oxford University Press.
- Darwall, Stephen (1983). *Impartial Reason*. Ithaca: Cornell University Press.
- Gensler, Harry (1985). "Ethical Consistency Principles." *Philosophical Quarterly*. 35(139): 156–170.
- Greenspan, Patricia (1975). "Conditional Oughts and Hypothetical Imperatives." *The Journal of Philosophy*. 72(10): 259–276.
- Hampton, Jean (1998). *The Authority of Reason*. Cambridge: Cambridge University Press.
- Hare, R.M (1971). "Wanting: Some Pitfalls" in *Agent, Action, and Reason*. Binkley, Bronaugh, and Marras, (eds.), Toronto: University of Toronto Press, 81–127.
- Hill, Thomas (1973). "The Hypothetical Imperative." *The Philosophical Review*. 82(4): 429–450.
- Kant, Immanuel (1997). *Groundwork for the Metaphysics of Morals*. Mary Gregor, trans. Cambridge: Cambridge University Press.
- Korsgaard, Christine (1997). "The Normativity of Instrumental Reason." In Cullity and Gaut, eds., *Ethics and Practical Reason*. Oxford: Oxford University Press.
- Lehrer, Keith, and Thomas Paxson, Jr. (1969). "Knowledge: Undeclared Justified True Belief." *The Journal of Philosophy* 66(8): 225–237.
- Scanlon, T.M. (1998). *What We Owe to Each Other*. Cambridge: Harvard University Press.
- Schroeder, Mark (2004). *Slaves of the Passions*. Doctoral Dissertation, Princeton University.
- (forthcoming). "The Hypothetical Imperative?" To appear in the *Australian Journal of Philosophy*.
- (ms1). "Reasons and Agent-Neutrality."
- (ms2). "Weighting for a Plausible 'Humean' Theory of Reasons."
- Sidgwick, Henry (1981). *The Methods of Ethics*. Indianapolis: Hackett.
- Smith, Michael (1994). *The Moral Problem*. Oxford: Basil Blackwell.
- Wallace, Jay (2001). "Normativity, Commitment, and Instrumental Reason." *Philosophers' Imprint* 1(3): <http://www.philosophersimprint.org/001003>.
- Watson, Gary (1975). "Free Agency." *The Journal of Philosophy*. 72(8): 205–220.