Direct Moral Grounding and the Legal Model of Moral Normativity

Benjamin Sachs

Abstract: Whereas most moral philosophers believe that the facts as to what we're morally required to do are grounded by the facts about our moral reasons, which in turn are grounded by non-normative facts, I propose that moral requirements are *directly* grounded by non-normative facts. This isn't, however, to say that there is no place in the picture for moral reasons. Moral reasons exist, and they're grounded by moral requirements.

Arguing for this picture of the moral sphere requires playing both offense and defense; this article provides the defense. I defend this view against the objections (1) that it must deny that one is generally blameworthy for having violated a moral requirement, (2) that it implies the existence of genuine moral dilemmas, (3) that it runs counter to an obviously true view of how moral deliberation should work, and (4) that it cannot explain why it *feels* as though figuring about what one is morally required to do often takes the form of thinking about what one's moral reasons are.

Keywords: moral reasons, moral requirements, grounding, blameworthiness, moral deliberation, reasons-first

I. A New Understanding of the Role Moral Reasons Play

The basic idea of the reasons-first program in practical philosophy, as recently developed by such philosophers as Michael Smith, T.M. Scanlon, Jean Hampton, Joseph Raz, Derek Parfit and John Skorupski, is to make the concept of a reason foundational in normativity and to ground all of the other key normative concepts—concepts from morality, prudential rationality, and perhaps even epistemology—in terms of a reason. The payoff would be an improved understanding of the other concepts themselves—including concepts of desire, value, and various deontic concepts—plus an explanation of why reasons are related in patterned and predictable ways to desires, values, and deontic facts. In addition, the boldest versions of the

program promise to metaphysically reduce the other normative entities to reasons, thereby yielding a simpler, less extravagant conceptual toolkit for practical philosophy.

Among those in the realist or quasi-realist camp (my only intended interlocutors here), the least controversial element of the reasons-first program is the claim that oughts, requirements, permissions, etc. are grounded in reasons. This is the claim that all reasons-first theorists accept, while they may disagree over (for instance) whether reasons are *conceptually* foundational and therefore the other normative concepts are to be understood in terms of reasons and also about whether we should "pass the buck" to reasons when it comes to desire and/or value.

Applied specifically to the moral domain, the claim would be that *moral* oughts, requirements and prohibitions (what I call 'verdictive facts') are grounded in *moral* reasons. Philosophers from across the spectrum of normative ethics believe this, including act-consequentialists (e.g. Kagan (1989: 65; 1988:17-8)), rule-consequentialists (e.g. Hooker (2002)), non-consequentialists (e.g. Nagel (1986:203), Baier (1958: ch. 3), Urmson (1974), Rachels (1999:196)), and even people who reject both consequentialism *and* non-consequentialism—particularists (e.g. Dancy (2004)). Of course they, like nearly everyone, accept some version of the claim that moral facts supervene on non-normative facts. Therefore, assuming that moral reasons don't reduce to non-normative facts, they must claim in addition that moral reasons are grounded in non-normative facts. So their picture of the moral domain is

> **Indirect Moral Grounding:** All the verdictive facts are grounded by moral reasons facts, which are all grounded by non-normative facts.

I, on the other hand accept:

Direct Moral Grounding: All the verdictive facts are grounded by non-normative facts.

One immediate question that will be asked of anyone who, like me, accepts Direct Moral Grounding (DMG), is whether there are moral reasons and, if so, how are they grounded. My answer is that moral reasons work in much the same way that, on a plausible theory of legal normativity, legal reasons work:

On this view, one's legal requirements ground one's legal reasons, not the other way around.[1] In other words, to find out what one has a legal reason to do, one looks to the laws to discover what one is legally required to do. Granted, there are often non-legal reasons to do the things that there are legal reasons to do. So, for instance, there is a prudential reason for me to wear my seatbelt, as the law requires. But there are no legal reasons aside from the ones grounded in legal requirements. In other words, actions obtain their legal significance by playing a certain role in the legal system.

This model has two important elements. First, legal requirements ground legal reasons. Second, there are no legal reasons other than those grounded in legal requirements. I now wish to add a third element to the model: the claim that all legal requirement-grounded reasons are of the same strength. The basic, and to my ear

---

[1] Although I do not intend to defend this theory of legal normativity here, I'd be remiss if I didn't take this opportunity to defuse one likely objection to it. This theory of legal normativity, like DMG, is a claim about grounding. It is *not* a conditional. If it were a conditional there would be serious trouble for it. It would be the claim: if one is legally required to φ then one has a legal reason to φ. This claim is subject to all sorts of good counterexamples having to do with unjust laws. This theory of legal normativity, since it is a grounding claim, merely says that a legal reason to do something is always grounded by a legal requirement to do that thing.

plausible-sounding, idea of this third element is that the fact that "it's the law" counts equally strongly in favor of each action it favors.[2]

I now want to propose a Legal Model of Moral Normativity (LMMN), which is basically the claim that what is true of legal reasons, on the proposal just made, is true of moral reasons as well. So LMMN is the claim that

- moral requirements ground moral reasons, and

- there are no moral reasons other than those grounded in moral requirements (though there are often reasons of other kinds to do the things that one is morally required to do), and

- all moral reasons are of the same strength.

I accept the conjunction of DMG and the LMNN mainly for two reasons: that there are no strong objections to it and that there *are* strong objections Indirect Moral Grounding (IMG).

My aim here is mostly to defend DMG/LMMN against the most likely objections against it. In each of the remaining sections I take up a single objection to DMG/LMMN and attempt to diffuse it. The other part of my case—identifying the defects in IMG—will have to be left aside because of space constraints, though it can be found elsewhere.[3]

---

[2] One might wonder why, then, it is appropriate for some violations of the law to be punished more severely than others. (Thanks to John Broome for alerting me to this question.) It wouldn't be appropriate, I admit, if the sole legitimate purpose of criminal punishment were a legal purpose—namely upholding respect for the law as such. And I concede that punishment requires a purely legal *justification* (i.e., punishment is justified only in cases of law-breaking). But justification is merely about when it's OK to punish; it leaves aside the separate question of *how* to punish when it's OK to punish. To answer that latter question we need to inquire as to the legitimate purpose(s) of punishment, and we may find that punishment has a legitimate non-legal purpose—for instance, promoting the social good, in which case we would have an argument for the legitimacy of using punishment for deterrence and prevention. And the deterrence and prevention motives could explain why it's appropriate for some violations of the law to be punished more severely than others—*viz*, some violations of the law are more important to deter and prevent than others.

[3] Benjamin Sachs, *Explaining Right and Wrong* [unpublished manuscript], ch. 3.

2. The Blameworthiness Objection

Consider the following argument:[4]

>P1. If X is morally required to φ, then X would be blameworthy for failing to φ.

>P2. If X would be blameworthy for failing to φ, then the balance of X's reasons favor X's φ-ing.

>C. Therefore, if X is morally required to φ, then the balance of X's reasons favor X's φ-ing.

This is a valid argument with compelling premises, and LMMN is inconsistent with its conclusion. This is the first objection to DMG/LMMN.

If we accept LMMN, then whenever one is morally required to do something one either has only a moral reason to do that thing or one has a moral reasons to do something else as well, but the two reasons are equally strong. The first sort of case is compatible with the conclusion of the above argument, but the second sort of case is not. To allow for the second sort of case, I propose substituting for P1

>P1*: If X is morally required to φ, then X will be blameworthy if she fails to φ.

The difference between P1 and P1* is that P1 does, while P1* does not, insist that if X is morally required to φ then if she fails to φ she must be blameworthy *for* failing to φ. P1* merely insists that X must be blameworthy. This allows her to be blameworthy for getting herself into a situation in which φ-ing was morally required

---

[4] Arguments quite like this one can be found in Darwall (2006: 28) and Portmore (2011: 43-4); see also their citations of other similar arguments.

of her, as opposed to being blameworthy for what she did once she found herself in it. This, I think, allows us to say what we want to say about cases of the second sort, just discussed (assuming there are such cases)—cases in which one has two moral requirement-generated reasons of the same strength: we want to say you can't be blamed for abiding by one requirement and not the other. But we want to say that you can be blamed for getting yourself into that situation. (Then again, maybe it's possible to blamelessly end up in such a situation. If so, we would want to reject P1*. But we would also be robbed of our grounds for supporting P1. Obviously I would be content with this turn of events.)

If we replace P1 with P1*, then the above argument becomes invalid. Yet P1* still gets us a strong connection between moral requirements and blameworthiness. As to the existence of a strong connection between moral requirements and moral reasons, which is what the above argument was supposed to establish, we can have that too, on LMMN, because LMMN implies that morality never requires one to act contrary to one's moral reasons.[5] So although accepting LMMN forces us to dismiss as unsound the argument given above, we can have everything we want nonetheless.

3. The Genuine Moral Dilemmas Objection

Because DMG/LMMN doesn't rule out that there may be multiple kinds of non-normative fact that can ground verdictive facts and says (1) that all moral requirements ground moral reasons, (2) that there are no other moral reasons aside from the requirement-grounded ones, and (3) that all moral reasons are of equal strength, DMG/LMMN seems to allow that there are situations that take the following

---

[5] As mentioned above, given LMMN whenever one is morally required to do something one either has only a moral reason to do that thing or one has a moral reasons to do something else as well, but the two reasons are equally strong.

form: one is morally required to φ and has a corresponding reason of strength x to φ, and one is also morally required to not-φ and has a corresponding reason of strength x to not-φ. There are no other moral reasons either to φ or to not-φ, so the overall moral reasons in favor of these two courses of action are equally strong. So it seems that according to DMG/LMMN morality can give us two incompatible moral requirements and no further guidance on what to do. Such a situation is known as a *genuine moral dilemma*.

This is troubling. And in fact, avoiding a picture of the moral domain on which moral conflicts have this shape constitutes a strong motivation—perhaps the quintessential motivation—for accepting IMG instead of DMG. As Broome (1999:412) puts it, "Weighing is just what reasons are made for."

There are well-known problems for theories that allow genuine moral dilemmas, having to do with their apparent violation of ought-implies-can and other intuitively plausible tenets of deontic logic (see Brink 1994). But I am going to set aside these concerns, for two reasons. First, they're not particular to DMG/LMMN. Some versions of IMG will allow for genuine moral dilemmas—namely, those that allow for the existence of trumping reasons. (For instance, if there is (as one might plausibly think there is) a trumping reason against choosing to end the life of one's child, then in *Sophie's Choice* Sophie faced a genuine moral dilemma.) Second, it would be an enormous undertaking to try to refute all the arguments against the possibility of genuine moral dilemmas, and I don't consider myself up to the task.[6]

So I'm going to proceed as if the problem for DMG/LMMN is simply that it's inconsistent with the widespread belief that genuine moral dilemmas are *rare*. Many

---

[6] But I would like to point out that I've already put a dent in the ought-implies-can argument against the possibility of genuine moral dilemmas. Insofar as there is an argument that ought implies can, it's likely to proceed by way of blameworthiness. The idea would be that ought-but-failed-to implies blameable-for, and blameable-for implies could have. But I've already argued, in §2, against the first premise of this argument.

would allow, I suspect, that some of what we would, in ordinary discourse, refer to as 'tragic' dilemmas are genuine moral dilemmas, but at the same time deny that there are genuine moral dilemmas outside tragic contexts. DMG/LMMN has the opposite implication on the (reasonable) assumption that there are multiple kinds of non-normative fact that can ground verdictive facts and that some of these non-normative facts are frequently instantiated.

I have two things to say by way of response to this objection to DMG/LMMN.

3.1 *The First Response to the Objection*

First, this counterintuitive implication of DMG/LMMN comes hand-in-hand with a theoretical virtue: it allows us to say what we want to say about moral residue. Most of us want to say that oftentimes, even when one does as morality requires, one's action leaves a moral residue, which means that it is appropriate to feel regret or guilt over what one has done and/or that one is morally required to make restitution to someone on account of what one has done. So, for instance, even if the promise I break is trivial, and I do it in order to avoid telling a big lie, we would think that I would owe something to the person I broke the promise to. DMG/LMMN has a wonderful explanation for this: sometimes even when one does as morality requires one also in so doing violates a moral requirement, because the situation one faced was a genuine moral dilemma. Surely it is appropriate to feel regret or guilt any time one violates a moral requirement, and surely one owes restitution to an individual with respect to whom one has violated a moral requirement.[7]

Of course the ability of DMG/LMMN to explain moral residue wouldn't count for much if moral residue could be accounted for in some other way. And in fact

---

[7] Sinnott-Armstrong (1988) may have been the first philosopher to use moral residue as an argument for genuine moral dilemmas.

there are several philosophers who hold that moral residue can be explained from within the framework of IMG. Since IMG is the view that a moral requirement to φ is grounded in the moral reasons for and against φ-ing, the obvious move for the supporter of IMG to make is to suggest that moral residue is explained by the fact that sometimes when the balance of moral reasons favors φ-ing there are also moral reasons favoring not-φ-ing, and consequently one often acts contrary to moral reasons even when doing what one is morally required to do. Oberdiek (2004:333), for instance, makes this move.

The problem, however, is that unfulfilled moral reasons don't seem to be per se the kind of thing that can give rise to moral residue. If it's appropriate for me to feel guilty over what I've done, or if I owe restitution to someone for what I've done, it seems it would have to be the case that somehow I failed to do what I *should* have done. It's not enough that my action wasn't perfect; that it left something to be desired. Instead, it should be that something was legitimately asked of me, and I failed to do it. But reasons (and this includes moral reasons) do not per se make demands of us; they're not deontic. It's merely in their nature per se to favor an action. To bolster this point, remember that if we are operating under the belief that moral reasons are directly grounded in non-normative facts, as IMG says, then it seems we should be quite liberal about them; we should admit that moral reasons come cheap. We should admit, for instance, that the fact that φ-ing would help someone grounds a moral reason to do it. But surely a failure to act on such a reason isn't necessarily something that it's appropriate to feel guilty about, or even something to regret, for that matter. If I pass up the opportunity to tell someone a joke (as I do many times a day, every day), and doing so would've brightened that

person's day, it would be odd if I came to regret that (barring the existence of some other consideration in favor of doing it).

The response will have to be that only *some* moral reasons can, when unfulfilled, give rise to moral residue. This suggestion is similar to what has been, historically, the most popular way to explain moral residue: positing that sometimes when we do as we are morally required to do we nevertheless violate a pro tanto obligation or duty.[8] What the two suggestions have in common is the idea that within the class of morally relevant non-normative facts, only some are such as to ground morally normative entities that when unfulfilled can give rise to moral residue, namely pro tanto obligations or duties or moral reasons of a certain kind.

To set up the discussion, let's first ask whether there are such things as pro tanto obligations or duties. Not even Prichard or Ross thought there were. Prichard spoke of "obligatoriness" but not of "obligations". Meanwhile, Ross said that a pro tanto duty, or a 'prima facie duty' as he called it, isn't a thing but rather "the characteristic…which an act has, in virtue of being of a certain kind (e.g. the keeping of a promise), of being an act which would be a duty proper if it were not at the same time of another kind which is morally significant" (1988:19). In other words, if an action has a characteristic that would make that action morally required (or 'duty proper', as Ross would say) in the absence of any other relevant considerations, we call that characteristic a pro tanto duty. But this sounds exactly like what a present-day adherent of IMG would say a moral reason is![9] So the strategy of explaining moral residue by reference to pro tanto duties, which Ross himself adopted (1988:28),

---

[8] Hare, (1981: chap. 2), Conee (1987: 241), McConnell (1996: 39-44), Brink (1994: 220-3), and Dahl (1996: 94).

[9] David Brink (1994) quite explicitly translates Rossian claims about pro tanto duty into the language of reasons.

is at risk of being indistinguishable from the proposal to explain it by reference to moral reasons.

To avoid this fate we need to part ways with Prichard and Ross and say that a pro tanto obligation/duty really is a thing. But what thing is it, and why should an adherent of IMG believe such things exist? She doesn't need pro tanto obligations or duties to fill the role of grounding the verdictive facts; she already has moral reasons at hand to fill that role. To posit that, in addition to or instead of moral reasons there are pro tanto obligations or duties, is to postulate a new entity simply for the sake of explaining moral residue.[10] So if we asked the IMG supporter what a pro tanto duty or obligation is, the only way she could say something helpful would be by saying (perhaps among other things) that pro tanto duties/obligations are what explain moral residue. Consequently, when the adherent of IMG tells us that moral residue arises in a certain case due to the presence of an unfulfilled pro tanto obligation or duty, she simply tells us that moral residue arises in that case due to the presence of something whose lack of fulfillment can give rise to moral residue. And this is no explanation at all.[11]

I conclude from all this that there are no resources within IMG to explain why moral residue, including the appropriateness of resentment and the demand for restitution, frequently arises. Given this, it really does matter that DMG/LMMN can

---

[10] Russ Shafer-Landau suggested to me that a pro tanto duty is an invariably relevant moral reason. In this case, postulating pro tanto duties doesn't amount to postulating a new entity. But by the same token it fails to provide a good explanation of moral residue, for the reason given earlier in this sub-section.

[11] A different strategy, proposed by Shpall (2014), is to explain moral residue by positing that sometimes when we do as we are morally required to do we nevertheless violate a commitment. This is in one way better than the pro tanto duty/obligation strategy, since we have an independent grip on what a commitment is (everyday normative talk includes talk of 'commitments'). But it is in another way worse, because it doesn't hold out the promise of explaining the whole range of cases in which moral residue arises. As Shpall more or less admits, becoming morally committed to something requires an act of will; I must make a promise, sign my name, take a vow, etc. Yet there are cases in which moral residue arises for some agent in spite of his not having engaged in any relevant act of will. Suppose, for instance, that I face a situation in which I must either harm A or harm B. If I harm A, I might owe restitution to her, despite the fact that I never engaged in any act of will that committed me to refraining from harming her.

11

do so by asserting that genuine moral dilemmas occur ore often than just in tragic

cases. This, of course, doesn't mean that there's nothing to be said against making

this assertion. It simply shows us that each normative ethical theory must find some

way of navigating the twin threats of countenancing frequent genuine moral dilemmas

and failing to explain the frequency of moral residue. My theory goes to the extreme

of entirely avoiding the latter problem while running smack into the former. But

every normative ethical theory must take some position or other along this spectrum,

and there's no position to be taken that doesn't have a cost.

3.2 *The Second Response to the Objection*

Second, strictly speaking DMG/LMMN says nothing about how frequent

moral dilemmas are. It is not a claim about *which* non-moral facts can directly

ground verdictive facts, and consequently it implies nothing about how frequent

moral conflict is. As for moral dilemmas, all it says is that however many moral

conflicts there are, every one of them is a genuine moral dilemma. Still, however, it

may yet turn out that DMG/LMMN succumbs to this objection. We will know this to

be the case if, in her effort to avoid countenancing a lot of moral dilemmas, the

DMG/LMMN-supporter is forced to adopt a version of the theory that is highly

implausible in some other way. To lay my cards on the table, my strategy for

avoiding this objection is to endorse a version of DMG/LMMN that says that there

are rather few facts about an action that can make that action impermissible or

required. This is something I actually want to say independently of the fact that it

helps me to avoid the worry about genuine moral dilemmas, and I argue elsewhere

that this is a strength in its own right.[12] But if I and any other DMG/LMMN-

---

[12] Sachs, *Explaining Right and Wrong*, chs. 4, 6-7.

supporters cannot be persuasive on this point, then, yes, DMG/LMMN is ultimately doomed by the worry about genuine moral dilemmas.

4. Deliberative Rationalism

The third objection to DMG/LMMN concerns moral deliberation. DMG holds that facts about moral reasons don't ground verdictive facts. If this were true then it would be hard to imagine that thinking about one's moral reasons could be a way of deliberating about what one morally ought to do. I expect this will strike many people as highly counterintuitive. Many will want to say, I suspect, that if you want to figure out what you morally ought to do then thinking about your moral reasons is precisely what you should do. I'll call this view *Deliberative Rationalism.* In this section I aim to offer an *independent* argument (an argument that doesn't use DMG as a premise) that Deliberative Rationalism is false.

4.1 *The Initial Argument and a Concession*

The problem for Deliberative Rationalism is that when it comes to determining the strengths of moral reasons, we have no clue what we're doing. And this, informally, is my objection to it. (I'll give my official version of the objection later.)

Now let me qualify what I just said. Urmson (1974:119) suggests we can appeal to intuition to weigh up reasons, and I admit that this is true. In really simple cases, such as a case where the choice is between killing an innocent person and breaking a trivial promise, we can generate intuitions about the comparative strengths

13

of the reasons at hand.[13] But in that case it would be pointless to use them. In such a case one will have a verdictive intuition in which one is confident on reflection, and this intuition will inevitably yield the same verdict as one would have reached by consulting one's intuition about the comparative strengths of the reasons and inferring a verdict on the basis of them. So one might as well just rely on one's verdictive intuition in determining what one is morally required to do.

By contrast, I submit that in hard cases (i.e., cases about which one does not have a verdictive intuition in which one is confident on reflection), where it would be quite helpful to have intuitions about the comparative strengths of the reasons at hand, we cannot generate such intuitions. Since this is a psychological claim and I'm not aware of any empirical studies on the matter, all I can do to support the claim is invite the reader to introspect:

Imagine you faced a situation in which you had promised a friend you would pick her up in the airport. However on your way there you encounter some construction, which you knew about but had forgotten to factor in to your decision about when to set off for the airport. If you follow the detour, you'll be very late. Your other option is to drive through the blocked-off zone, thus endangering and frightening the construction workers on-site. Now ask yourself, Which moral reason

---

[13] Scanlon (2014: ch. 5) points out that in certain cases there are other methods available to us besides appeal to intuition. First, there are what Scanlon calls 'quantitative cases,' where the strengths of the relevant reasons correspond to something quantitative. This would be the case, Scanlon says, when one is trying to decide which of two cars to buy and the consideration in question is the matter of price. Scanlon goes on to point out that quantitative cases (where the relevant quantities are unequal) constitute a subclass of dominance cases. Dominance cases are another kind in which we can determine all we need to know about how strong the relevant reasons are. If I can achieve X by φ-ing, but can achieve X and Y (another good thing) by ψ-ing, then obviously the reason in favor of ψ-ing is stronger than the reason in favor of φ-ing. The final kind of case Scanlon mentions is the 'multi-level case'. This is where one has a higher-level reason to place greater weights on one kind of reason in preference to another. So if one is on a diet one has a higher-level reason to place greater weight on the reason favoring following one's diet than that placed on one's reason to deviate for the sake of pleasure or convenience.

  What's important, for our purposes, is that Scanlon's methods are available only for simple cases (aside from the odd case involving higher-level reasons of a particular sort). But as I'm about to argue, the real problem for Deliberative Rationalism is the hard cases. (In fairness to Scanlon, he does not set out to demonstrate that we can determine the relative strengths of reasons in hard cases.)

is stronger: The moral reason favoring following the detour, grounded in the well-being of the construction workers, or the moral reason favoring ignoring the detour, grounded in your promissory obligation to your friend?

My conjecture (based on my own introspection) is that when we ask ourselves questions like this we answer them inferentially. First we ask ourselves what would be the right thing to do in the situation, and then we extrapolate from that an answer to the reason-strengths question. (Of course a procedure as simple as this wouldn't be available for any case in which more than two moral reasons were involved. More on this later.) I presume the reader's own introspection yielded the same result.

So here's the situation: In simple cases one's intuition about the relative strengths of the competing reasons will always mirror one's verdictive intuition, and so one might as well just appeal directly to one's verdictive intuition. And in the hard cases we don't have intuitions about the relative strengths of the competing reasons. To put this colloquially: In hard cases, if the right thing to do comes down to how the reasons on either side tot up, then there had better be quantities to tot. But there never are, for all our intuitions can tell us.[14]

---

[14] I first raised this worry specifically against Douglas Portmore's consequentialism. See Sachs (2013). A similar point has been made by Richardson (1990: 287-9). McMahan (2000:105) confronts the same problem and simply bites the bullet: he says that the direction of discovery is the reverse of the direction of grounding. What this means, I think, is that he accepts IMG but rejects Deliberative Rationalism.

So far I have mentioned just two ways of arriving at a reason-strength belief: (1) having a reason-strength intuition, and (2) inferring such a belief from one's verdictive beliefs. I don't want to deny that there might be other ways of arriving at a reason-strength belief; I simply maintain that the other methods don't solve the problem for Deliberative Rationalism. Here are three case studies:

Kearns and Star believe that a reason to φ is evidence that one ought to φ, and the strength of any reason is the strength of the evidence it constitutes, where the strength of a piece of evidence is the degree to which it raises the probability of whatever proposition it is evidence for (Kearns and Star 2009: 230-2; see Thomson 2008 for a similar view). For our purposes the question is: If Kearns and Star were right about what a reason is and what determines its strength, would this make available a new defense of Deliberative Rationalism? The answer, unfortunately, is no. One will have no greater ease in directly discerning the degree to which some fact makes it more likely to be the case that one ought to φ than one would have in directly discerning the strength of some reason to φ.

A similar result emerges if we adopt thoroughgoing Humeanism, where 'Humeanism' is the view that each reason is grounded in a desire, 'thoroughgoing Humeanism' is a view, such as Michael Smith's, on which the strength of a reason is determined by the strength of the desire that grounds it (1996: 167). Here the question is: If Indirect Moral Pluralism were true and Smith was right about what a

4.2 The Uselessness of Inferential Beliefs about Reason-Strength

In order to show that thinking about one's moral reasons is a way of thinking about what one morally ought to do, one needs to show that one of the two allegations I made above—i.e., that thinking about one's moral reasons is useless in easy cases and that thinking about one's moral reasons is useless in hard cases—is false. I assume that the better strategy is to challenge the latter, since the former seems obviously true. What, then, is the best thing for a supporter of Deliberative Rationalism to say regarding deliberation about hard cases? What she would say, I assume, is that we can work in both directions. Specifically, when trying to determine what one morally ought to do in some hard case, A, we can

1) identify the moral reasons relevant to A: N_1, N_2, N_3, …, N_X, then

2) identify other situations—B, C, D, …, Z—in which the same reasons apply, then

3) appeal to intuition to determine what one ought to do in B, C, D, …, Z, then

---

reason is and what determines its strength, would this make available a new answer as to how one can determine how strong one's reasons are? Again, the answer is no. Since the relevant desires, for Humeans, are always hypothetical ones, we would have a good method for discerning the strengths of reasons, given thoroughgoing Humeanism, if and only if we had a good method for discerning the strengths of hypothetical desires. But, of course, we don't.

The final case study is Mark Schroeder's (2007) theory of reason-strengths. First, for clarification, his view is that reasons don't have weight, but rather sets of reasons do (126). And one of the ways for one set of reasons, A, to be weightier than another, B, "is for the set of all the (right kind of) reasons to place more weight on *A* to be weightier than the set of all (the right kind of) reasons to place more weight on *B* (138)." So if we're to have a reliable way of figuring out where the balance of reasons to act lay, we need a reliable way of figuring out where the balance of reasons to place weight on reasons lay. Do we? Schroeder says some helpful things (141-4), but they do not, nor do I think they're intended to, actually constitute a useable method for figuring out where the balance of reasons to place weight on reasons lay. Providing such a method isn't really necessary given the agenda of Schroeder's book. He's interested in showing how we can *explain* why there are weightier reasons for certain people in certain situations to do certain things—and, in particular, that a thesis that connects reasons to desires provides a good explanation—as opposed to showing how we can *figure out* which reasons for certain people in certain situations to do certain things are weightier than which.

4) use that information to determine the relative strengths of N_1, N_2, N_3, ..., N_X, then

5) use that information to determine what one morally ought to do in A.

Sometimes we know a verdictive fact and can use that knowledge to infer facts about the relative strengths of certain reasons. (This is precisely how Joshua Gert (2005) suggests we can discover the relative strengths of reasons.) Then we can take that knowledge and work in the other direction, inferring further verdictive facts from the just-discovered facts about the strengths of various reasons.

This is sound methodology, I admit. But it doesn't get us what we want. Moral deliberation isn't supposed to merely fill in the blank spots in our verdictive belief-set; it's supposed to provide independent access to the verdictive facts. Now just about everyone agrees that one way to improve one's verdictive belief-set is to discard or at least adopt a cautious stance toward those of one's verdictive intuitions that one recognizes were formed under circumstances not conducive to sound moral thinking; this bit of methodology can be accommodated by anyone. But the moral deliberator should want more; she should want a way of being shown that some of her verdictive beliefs—even the ones formed under ideal conditions—are false. And this is exactly what we deliberators lack if Deliberative Rationalism is true and we don't have intuitions about reason-strengths in hard cases. Without those intuitions, we don't have a second set of moral inputs in hard cases (as opposed to easy cases, where in addition to verdictive intuitions we have reason-strength intuitions).[15] Moral reasoning becomes a matter of figuring out how much information one can extract

---

[15] Actually, our metaethical intuitions may serve as a second set of inputs, but still this doesn't get us very far. If our metaethical intuitions provide any constraints at all on our verdictive beliefs, surely those constraints are merely formal. (For instance, we might have the metaethical intuition that like cases should be judged alike. If we take this intuition seriously then some pairs of verdictive beliefs are ruled out.) And formal constraints don't constitute much of an independent check on our verdictive intuitions.

from one's verdictive intuitions plus the non-normative facts. But if one reasons in this way then one will never be pressured to alter one's verdictive belief-set.[16] And that, officially, is my objection to Deliberative Rationalism.

4.3 *Turning the Tables*

Let me recap. If DMG is true and in morality the direction of discovery is the same as the direction of grounding, then one picture of moral deliberation that *cannot* be true is Deliberative Rationalism. This was supposed to be a reductio of DMG, because Deliberative Rationalism is so appealing initially. But I've shown that Deliberative Rationalism is false.

Now let's ask the natural follow-up question. Given that I criticized Deliberative Rationalism on the grounds that since we do not have beliefs about reason-strength in hard cases that are supportable independently of our verdictive beliefs we will never be forced to alter our verdictive beliefs about hard cases if we deliberate in the way that Deliberative Rationalism suggests, we should want to know whether the situation appears at all different given the denial of Deliberative Rationalism. Specifically, we should want to know whether the situation is different given whatever picture of moral deliberation follows from DMG plus the assumption that the direction of discovery is the same as the direction of grounding.

If the direction of discovery is the same as the direction of grounding, and (as is DMG would have it) non-normative facts ground the verdictive facts, then to discover the verdictive facts it is necessary to first discover the relevant non-normative facts. But this might not be sufficient, since saying that the non-normative facts ground the verdictive facts isn't equivalent to saying that the non-normative

---

[16] Herman raises a worry like this in "Kantian Commitments" (unpublished manuscript).

facts that *can* ground the verdictive facts always *do* ground them when they (the non-normative facts) are instantiated. And, indeed, it's most plausible to suppose that they don't always do so. For instance, although an utterance's being a lie is (plausibly) a non-normative fact about an utterance that can ground the impermissibility of uttering it, it doesn't seem as though the fact that an utterance is a lie *always* grounds the impermissibility of uttering it. An exception, perhaps, would be a case in which the lie is about the strength of the poker hand you're holding.

What we might say, to account for this, is that for some (or perhaps all) non-normative facts, n, that can ground verdictive facts, there is one or more fact the instantiation of which prevents n from grounding the verdictive fact it would otherwise ground. Borrowing from Jonathan Dancy (2004), I call these preventing facts *disablers*. I argue elsewhere that we should indeed concede the existence of disablers.[17]

Given the existence of disablers, discovering the verdictive fact about a given action requires, according to the DMG-supporter, identifying not only whether any non-normative facts that can do grounding work are instantiated but also whether any relevant disablers of those non-normative facts are instantiated.

Returning now to our question of whether DMG can offer an attractive alternative to Deliberative Rationalism, we seem to encounter a problem. If we admit the existence of disablers, it may appear that the DMG-supporter has an easy way to escape having to accept any verdicts she finds counterintuitive. If she accepts that non-normative facts of kind N can ground verdictive facts of kind V, and there is a situation in which a fact of kind N is instantiated but according to her intuitions no

---

[17] Sachs, *Explaining Right and Wrong*, ch. 1.

fact of kind V is instantiated, she can just claim that some other kind of fact, Z, which is also instantiated, is a disabler of facts of kind N.

This escape route, however, will often not be available to the DMG-supporter. We *can't* say everything we want to say about disablers. Some of the things we might want to say would be drastically counterintuitive. In other words, *given DMG we have the all-important second set of inputs to moral deliberation—our intuitions about disabling*. Suppose, for instance, A promises to help B with B's taxes, but then when the time comes A has another obligation that she needs to discharge at the same time, such that she cannot discharge both. Depending on what that other obligation is, we might intuitively want to say that A is morally required to discharge that other obligation and not morally required to help B with B's taxes. But to say this we'd have to claim that A's obligation to provide help with B's taxes is disabled. And that is drastically counterintuitive. It just doesn't seem as though having a conflicting obligation could disable the moral importance of keeping a promise. So we can't get away with saying what we want to say in this case.

The moral epistemology I sketch here occupies an attractive intermediate point between absolutism and minimalism. Absolutism is the view that normative ethics should be aimed at the discovery of general moral truths (however highly specified and complex) from which one can infer verdictive facts without the use of judgment. I am not an absolutist, as I hold that even once we have identified the general truths as to which non-normative facts can ground the verdictive facts we cannot apply this knowledge to cases without the use of judgment (specifically, judgment as to the presence or absence of disablers). Minimalism is the view that it's no mark against a normative ethical theory if it leaves such a wide scope for judgment that the theory will never deliver verdictive conclusions that we find counterintuitive.

(I have in mind here the epistemology implied by IMG and also by the Prichard/Ross-style view endorsed by McNaughton (2002).) I am not a minimalist, as I propose that we should consider ourselves constrained by our intuitions when exercising our judgment about the presence or absence of disablers in particular cases.

5. The Phenomenology Objection

I've argued that thinking about one's moral reasons is not a good method of moral deliberation. But many people report that when they think their way through tough moral cases it *seems* as though their deliberation is deliberation about how the moral reasons on each side balance out (Kagan 1988:19).[18] So I have some explaining to do.

I want to offer a three-part explanation. First, I should point out that nothing I've said here commits me to the view that it's pointless to deliberate about one's moral reasons. For while I've said that such deliberation can't help one to figure out what one morally ought to do, I *can* allow that it can help one to figure out what one ought to do all-things-considered. After all, I do admit the existence of moral reasons, and so I can admit that one's moral reasons help to determine what one ought to do all-things-considered. Now this explanation wouldn't get us very far if the phenomenology in question were the phenomenology of distinctly *moral*, as opposed to all-things-considered, deliberation. But I doubt that all-things-considered deliberation where moral matters are at stake would feel much different from purely moral deliberation. Importantly, in both cases one might feel anticipatory versions of

---

[18] Thanks to Miranda Fricker for pushing me to address this objection.

guilt, regret, etc., as one considers the option of failing to do what one is morally required to do.

Admittedly, there's one way that all-things-considered moral deliberation where moral matters are at stake could come to feel quite different from purely moral deliberation—namely if the person doing the deliberating had the belief that the considerations at stake either were or were not purely moral. This brings me to the second part of my explanation of the phenomenology: to the extent that we sometimes feel as though we are weighing moral reasons as a way of deliberating about what we morally ought to do, this may simply be because of false normative beliefs. There could be cases in which we balance multiple reasons against each other, where some of them are moral reasons and others are non-moral reasons, but we falsely believe the non-moral reasons to be moral reasons. To take an actually controversial case: some people believe (as W.D. Ross did) that if φ-ing would contribute to my self-perfection then I have a moral reason to do it. But maybe it's actually the case that this particular non-normative fact grounds a non-moral reason.

This point about false beliefs can be redeployed to explain other cases in which we deliberate about our moral reasons in order to determine what we morally ought to do. For instance, there could also be cases in which we balance multiple reasons against each other, where they are both moral reasons and have the same weight, yet we falsely believe that they have differing weights. After all not everyone accepts that part of LMMN that says that all reasons grounded in moral requirements are of the same strength. False beliefs like this make it *in fact possible* to deliberate about one's moral reasons as a way of arriving at a conclusion about what one morally ought to do.

The third part of my explanation involves an appeal to the distinction between practical deliberation and theoretical deliberation. Practical deliberation is deliberation about what to do, where the end of such deliberation is intending, or willing, or acting, or something like that. Moral deliberation is a kind of theoretical deliberation, since it is reasoning toward the fact about what one morally ought to do. (Though of course it's usually theoretical deliberation toward a practical end—the end of doing the right thing.) The key point is that all of the non-normative facts that would be relevant to moral deliberation are also relevant to practical deliberation. So if the phenomenology of moral deliberation in hard cases is the phenomenology of weighing non-normative facts thought of simply as facts (as opposed to non-normative facts thought of as reasons), then there might not be much difference in feel between moral deliberation about situation X and practical deliberation about situation X. In both cases, one's thought process could take the form: "On the one hand φ-ing would hurt someone, but on the other hand not-φ-ing would be deceptive…" There's nothing in DMG/LMMN that implies that practical deliberation of this sort is pointless, even in a genuine moral dilemma. The mere fact that all the moral reasons are in perfect balance does not imply that practical deliberation about the facts that ground those reasons cannot be productive. In practical deliberation we might legitimately take those facts into account in an entirely different way from how it would be appropriate to account for them in moral deliberation.

References

Baier K (1958) The moral point of view. Cornell University Press, Ithaca

Brink D (1994) Moral conflict and its structure. Philosophical Review 103:215-24

Broome J (1999) Normative requirements. Ratio 12:398-419

Conee E (1987) Against moral dilemmas. In: Gowans CW (ed) Moral dilemmas. Oxford University Press, New York, 239-49

Dahl NO (1996) Morality, moral dilemmas, and moral requirements. In: Mason HE (ed) Moral dilemmas and moral theory, Oxford University Press, New York, 86-101

Dancy J (2004) Ethics without principles. Clarendon Press, Oxford

Darwall S (2006). The second-person standpoint. The Harvard University Press, Cambridge

Gert J (2005) A functional role analysis of reasons. Philosophical Studies 124:353-78

Hare RM (1981) Moral thinking. Clarendon Press, Oxford

Hill TE (1996) Moral dilemmas, gaps, and residues. In: Mason HE (ed) Moral dilemmas and moral theory. Oxford University Press, New York, 167-98

Hooker B (2002) Intuitions and moral theorizing. In Stratton-Lake P (ed) Ethical intuitionism: re-evaluations. Clarendon Press, Oxford, 161-83

Kagan S (1988) Normative ethics. Westview Press, Boulder

Kagan S (1989) The limits of morality. Oxford University Press, New York

Kearns S, Star D (2009) Reasons as evidence. In Shafer-Landau R (ed) Oxford studies in metaethics, vol. 4. Oxford University Press, New York, 215-242

McConnell T (1996) Moral residue and dilemmas. In: Mason HE (ed) Moral dilemmas and moral theory. Oxford University Press, New York, 36-47

McMahan J (2000) Moral intuitions. In Lafollette H (ed) The blackwell guide to ethical theory. Blackwell Publishers, Malden, 92-110

McNaughton D (2002) An unconnected heap of duties? In Stratton-Lake P (ed) Ethical intuitionism: re-evaluations. Clarendon Press, Oxford, 76-91

Nagel T (1986) The view from nowhere. Oxford University Press, New York

Oberdiek J (2004) Lost in moral space: on the infringing/violating distinction and its place in the theory of rights. Law and philosophy 23:325-436

Portmore D (2011) Commonsense consequentialism. Oxford University Press, New York

Rachels J (1999) The elements of moral philosophy. Mcgraw-Hill, Boston

Richardson, HS (1990) Specifying norms as a way to resolve concrete ethical problems. Philosophy & Public Affairs 19:279-310

Ross WD (1988) The right and the good. Hackett, Indianapolis

Sachs B (2013) Reasons consequentialism. Journal of Moral Philosophy 10:671-82

Scanlon TM (2014) Being realistic about reasons. Oxford University Press, New York

Schroeder M (2007) Slaves of the passions. Oxford University Press, New York

Shpall S (2014) Moral and rational commitment. Philosophy and Phenomenological Research 88:146-72

Sinnott-Armstrong W (1988) Moral dilemmas. Basil Blackwell, Oxford

Smith M(1996) Normative reasons and full rationality: reply to swanton. Analysis 156:160-8

Thomson JJ (2008) Normativity. Open Court Publishing, Chicago

Urmson JO (1974) A defence of intuitionism. Proceedings of the Aristotelian Society 75:111-19