

Infinitism about Cross-Domain Conflict

David Killoren

7.1. INTRODUCTION

According to *the overriding view*, (i) there are multiple normative domains, including the domains of morality and prudence, and perhaps also including an epistemic domain, an esthetic domain, a domain of etiquette, and so on. (ii) Normative domains can come into conflict with one another. Such conflicts occur when an obligation from one domain is incompatible with an obligation from another domain. For example, an agent may be morally obligated to ϕ_1 and prudentially obligated to ϕ_2 in a case where the agent cannot do both ϕ_1 and ϕ_2 . And finally, (iii) in some such conflicts, an obligation belonging to one normative domain overrides an obligation belonging to another normative domain. For example, a moral obligation may override a prudential obligation or vice versa. In such cases, we may say, the first domain *trumps* the second domain.

My purpose in this paper isn't to defend the overriding view, but is instead to tease out some of its implications. So I won't say much in defense of the overriding view. I'll avoid discussion of the merits of (i) entirely. But I will say a bit about (ii) and (iii), before I explain the implications of the overriding view that will be my main focus in this paper.

There are theories that undermine (ii). For example, ethical egoism twists morality into something resembling prudence; therefore ethical egoists will not see any conflict between morality and prudence and thus will reject (ii) (Copp 1997: 89). Similarly, some moral theories, including certain kinds of Aristotelian eudaimonism, hold that it is in one's own self-interest always and everywhere to be morally excellent. Such theories twist prudence into something resembling morality, and so defenders of such theories will naturally resist (ii). And, as Parfit (2011: 142) observes, certain religious ideas rule out (ii) by positing a divine force that rewards right action and thus ensures that it's always prudent to do what's moral.

But, setting religious ideas aside for the present purpose, I think all of the theories that lead to the denial of (ii) are implausible, precisely because they erase a distinction and a conflict that is an obvious and pervasive part of everyday normative experience. Morality seems very often to require us to behave imprudently. For example, we seem often to be morally obliged to make ourselves worse off (or to forego benefits for ourselves) in order to help (or to avoid doing harm to) those who are disadvantaged, oppressed, or helpless.

The conflict between morality and prudence is a version of what Sidgwick (1874) calls “the profoundest problem” in moral philosophy. I don’t intend to solve the problem posed by such conflict in this paper, but I will offer a sketch of what I believe a solution ought to look like. My assumption, first, is that we would have a solution to the conflict between morality and prudence—or between the epistemic domain and the esthetic domain, or between any two normative domains one wants to consider—if we could show that (iii) is true. That is, my assumption is that a conflict between (say) morality and prudence in a given case would be resolved if we could show that morality wins out over prudence, or that prudence wins out over morality.¹

So my assumption is that *if* the overriding view, including (iii), were true, *then* we would have a solution to the problem posed by conflicting normative domains. My aim in this paper is to explore what follows from the overriding view. The thesis of this paper is: *If* we endorse the overriding view, *then* we ought to endorse the view that I call *infinitism about cross-domain conflict*.

Infinitism about cross-domain conflict is a direct analog of epistemic infinitism, a view that has been given its seminal defense by Peter Klein (1998). Epistemic infinitism is an alternative to epistemic coherentism and to epistemic foundationalism. Roughly, according to epistemic infinitism, a given belief that *p* is epistemically justified iff there is an infinitely long chain of reasons *q*, *r*, *s*, . . . , such that *q* supports *p*, *r* supports *q*, *s* supports *r*, and

¹ Sidgwick seems to have thought that a *satisfactory* solution to the present problem requires that morality wins and prudence loses; see Crisp (ms.). But many philosophers, e.g., Susan Wolf (2015), are sanguine about the possibility that prudence sometimes wins out over morality. In any case, this issue need not concern us here; an unsatisfying solution to a given problem is still a solution to that problem. Cf. McPherson (2018). A related point is that (iii) isn’t the *only* possible solution to the present problem. For instance, David Copp (ms.) argues for a kind of pluralism according to which, in the event of conflicts between morality and prudence or between any other pair of domains, there is simply no winner and no loser; see also Tiffany (2007) and Enoch (2019). If such a position can be given a good defense—a matter I won’t address in this paper—then it would count as a solution of sorts to the problem here.

so on, indefinitely. Infitism about cross-domain conflict can be defined analogously, as follows:

Infitism about cross-domain conflict

There is at least one case in which an agent has an obligation O_A that belongs to domain D_A , and an obligation O_B that belongs to domain D_B (where D_A and D_B are different domains), and O_A overrides O_B . For every such case, there is an infinitely long chain of domains D_1, D_2, D_3, \dots , such that D_1 contains an obligation to act in line with the obligations of D_A , and D_2 contains an obligation to act in line with the obligations of D_1 , and D_3 contains an obligation to act in line with the obligations of D_2 , and so on, indefinitely.

There is a deep literature on epistemic infinitism, but no one has given a careful argument for the view that I'm calling infinitism about cross-domain conflict. When this sort of view is mentioned, it is usually treated with casual disregard; e.g., Dale Dorsey (2016: 23) considers a view in the neighborhood and simply dismisses it as "a wildly unattractive picture of the normative enterprise."

My argument from the overriding view for infinitism about cross-domain conflict is a regress argument and as such it formally resembles many other regress arguments that have appeared in philosophy. The argument has an important precedent in Copp (1997). It is fruitful to compare similarities and differences between my argument and Copp's argument but unfortunately I do not have the space to provide that comparison here. My argument is as follows.

The argument for infinitism about cross-domain conflict

- (1) O_A overrides O_B iff O_A belongs to D_A and O_B belongs to D_B and D_A trumps D_B . (This is just a stipulative definition of trumping.)
- (2) If D_A trumps D_B , this is because there is a novel higher-order domain, D_1 , that explains why D_A trumps D_B .
- (3) D_1 explains why D_A trumps D_B only if D_1 has the requisite authority.
- (4) D_1 has the requisite authority only if D_1 trumps at least some of its same-order competitors (call them $D_1^*, D_1^{**}, D_1^{***}, \dots$).
- (5) If D_1 trumps any of its same-order competitors, this is because there is a novel higher-order domain, D_2 , that explains why D_1 trumps those same-order competitors.

Step (5) marks the start of a regress. From this point, the argument proceeds to a step analogous with (3), and then to a step analogous to (4), and then to a step analogous with (5), and so on, continuing in this way

indefinitely. This, as I'll explain in due course, leads to infinitism about cross-domain conflict.

Section 7.2 will develop the argument. Section 7.2.1 will lay groundwork for the argument; Section 7.2.2 will argue for each of the premises. Section 7.3 will address objections.

7.2. THE ARGUMENT FOR INFINITISM ABOUT CROSS-DOMAIN CONFLICT

In this section I will assume the overriding view and argue for infinitism about cross-domain conflict. In Section 7.2.1, I will argue for position concerning what I'll call *novel higher-order domains*. Specifically, I will argue that novel higher-order domains exist and that standpoints associated with them play a crucial epistemic role. Then, in Section 7.2.2, I will complete the argument for infinitism about cross-domain conflict. First I need to make some preliminary stipulations and definitions.

I'll first explain the notion of a normative standpoint. Suppose Monica has a moral reason to donate a sum to charity. In virtue of this reason, suppose, Monica is (all-things-considered) morally obligated to donate to charity. Further, Monica has a prudential reason to invest the sum in her retirement plan. In virtue of this reason, Monica is (all-things-considered) prudentially obligated to invest. Monica can't do both. Finally, suppose that it intuitively appears that Monica's moral obligation to donate to charity overrides her prudential obligation to invest. Then, in this case, it appears that cross-domain overriding occurs: Monica's moral obligation overrides her prudential obligation; morality trumps prudence.

From the moral standpoint, it seems Monica ought to donate, whereas from the prudential standpoint, it seems Monica ought to invest. Observing this, we can reasonably suppose that a normative standpoint is, broadly, a way of seeing or thinking about situations in which an agent must decide among options. This way of seeing or thinking might involve the (conscious or otherwise) application of general principles (e.g., general moral principles, or general principles of prudence) or it might involve something like intuitions about cases, or apprehension of substantive conceptual truths like "moral fixed points" (Cuneo and Shafer-Landau 2014). For my purpose here I want to be as neutral as possible about how precisely standpoints work psychologically.

What is important, for my purpose, is simply that we seem typically capable, when presented with any given situation in which multiple options are available, of viewing that situation from several different standpoints—such as the moral standpoint, the prudential standpoint, and so on. But what

makes it the case that this or that standpoint counts as the (or a) moral (or prudential, or aesthetic, etc.) standpoint? I think we can answer that question in terms of domains: A given standpoint is the moral (or e.g., prudential) standpoint if and because obligations in the moral (or e.g., prudential) domain are apprehensible from it.² In other words, I'm suggesting that a given standpoint counts as the X-standpoint because obligations of the X-domain are apprehensible from that standpoint.³

Let's say that a *zeroth-order* standpoint is a standpoint of the sort that most of us are already familiar with: a standpoint from which we can consider the question of what an agent is to do in a given situation. A *first-order standpoint*, then, is a standpoint from which we can consider the question of which zeroth-order standpoint the agent is to adopt and act in accordance with—or, equivalently, the question of which zeroth-order domain ought to guide the agent in deciding what to do. A second-order standpoint is a standpoint from which we can consider the question of which first-order standpoint the agent ought to adopt. A third-order standpoint is a standpoint from which we can consider the question of which second-order standpoint the agent ought to adopt. And so on. Let's say that a *higher-order standpoint* is any standpoint that lies above the level of zeroth-order standpoints.

So now we have an idea of what a higher-order standpoint is. A *novel* first-order standpoint is a standpoint that is distinct from any of the zeroth-order standpoints under consideration in a given case. Not all higher-order standpoints need be novel in this sense. The moral standpoint, for example, seems to double as both a zeroth-order standpoint and as a first-order standpoint: The moral standpoint is a standpoint from which we can consider questions about what an agent is to do *and* can consider questions

² Suppose that, in a given case, one's moral obligation is simply to do what one is prudentially obligated to do. Then you might think that in order to apprehend one's moral obligation from the moral standpoint, one must be able to apprehend one's prudential obligation from within that standpoint. But then the moral standpoint would (by definition) be the prudential standpoint. (Thanks to an anonymous referee for pushing me on a version of this point.) In order to avoid this result, I can say that in this sort of case, in order to apprehend one's moral obligation from the moral standpoint, one must first already have apprehended what one's prudential obligation is; and in order to know this, one must have previously occupied the prudential standpoint.

³ This defines *what it is for a standpoint to be the X-standpoint*; it does not define the notion of a standpoint itself. I take standpoints to be a familiar feature of normative experience and so I don't think it is necessary to define the notion of a standpoint, though it is important that standpoints characteristically allow us to apprehend obligations. Other features of standpoints might also be characteristic of them; e.g., it may be that standpoints are characteristically associated with certain types of reactive attitudes (cf. McElwee 2017).

about which first-order standpoint an agent is to adopt. The moral standpoint may—or may not—be self-ratifying at the first-order level; that is, it may or may not be that from the moral standpoint it looks as if any given agent ought to rely upon the moral standpoint in order to decide what to do. For example, if something like the two-level utilitarianism described by Hare (1981) were true, then morality would not be self-ratifying at the first-order level.

Throughout this paper I will assume *deontic rationalism*, the view that obligations of any given domain are explained entirely by reasons from that domain. So, for example, I am assuming that moral obligations are explained entirely by moral reasons, that prudential obligations are explained entirely by prudential reasons, and so on. Further, I will assume that to be a reason of any given domain is precisely to be capable of entering into an explanation of an obligation from that domain. So, for example, to be a moral reason (as against some other sort of reason) is just to be capable of explaining moral obligations.⁴

Deontic rationalism thus conceived implies that *at least some* reasons are in the business of explaining obligations (so, it's in tension with certain views, such as that of Kearns and Star 2008). But deontic rationalism does not say that *all* reasons are in the business of explaining obligations (so, it is consistent with the possibility of what Dancy calls 'enticing reasons'—reasons that favor but don't require). Deontic rationalism says that *if* a reason belongs to a given domain then it explains obligations in that domain.

7.2.1. From the Overriding View to Novel Higher-Order Domains

Consider the following pair of views:

Cross-domain reasons fundamentalism—metaphysical component:

Suppose obligation O_A belongs to normative domain D_A and obligation O_B belongs to normative domain D_B (where D_A and D_B are different domains). If O_A overrides O_B , this is *entirely because* the reason(s) that give rise to O_A outweigh the reason(s) that give rise to O_B .

Cross-domain reasons fundamentalism—epistemological component:

Suppose obligation O_A belongs to normative domain D_A and obligation O_B belongs to normative domain D_B (where D_A and D_B are

⁴ This does not entail that nothing is supererogatory. The assumption that all reasons are *capable* of explaining moral obligations does not imply that reasons are not also capable of explaining other sorts of properties, such as supererogation.

different domains). If O_A overrides O_B , this *can be known simply by observing (i.e., coming to know) that* the reason(s) that give rise to O_A outweigh the reason(s) that give rise to O_B .

To make the argument that I want to make in this section, I will first need to argue that the epistemological component of cross-domain reasons fundamentalism is false. So consider three possibilities:

INCOMPARABILITY: Reasons of one domain never outweigh or are outweighed by reasons of another domain. Reasons of different domains are, in short, incomparable.

COMPARABILITY WITHOUT DETERMINATION: There is cross-domain outweighing at the level of reasons—reasons of different domains are comparable—but cross-domain outweighing at the level of reasons cannot fully determine or explain cross-domain overriding at the level of obligations.

COMPARABILITY WITH DETERMINATION: There is cross-domain outweighing at the level of reasons, and cross-domain outweighing at the level of reasons fully determines and explains cross-domain overriding at the level of obligations.

I will argue that the epistemological component of cross-domain reasons fundamentalism is false regardless of which of the above theses is true.

7.2.1.1. *Incomparability and Comparability Without Determination*

If *Incomparability* is true, the weight of Monica's moral reason to donate can't be compared with the weight of Monica's prudential reason to invest. In that event there are two possibilities. The first possibility is that (I) there is no cross-domain overriding: so, e.g., Monica's moral obligation neither overrides nor is overridden by her prudential obligation. But we have to reject this possibility because, in assuming the overriding view, we've assumed that there *is* cross-domain overriding. The second possibility—which is therefore the only one open to us here—is that (II) when a given obligation O_A overrides another obligation O_B (where O_A and O_B belong to different domains) this is explained in terms of something other than the relative weights of the reasons underlying O_A and O_B . For example, perhaps there is something special about morality (or about prudence, or whatever) in virtue of which moral obligations (or prudential obligations, or whatever) override other types of obligations, regardless of the relative weights underlying the obligations in question.

If (II) holds, then cross-domain reasons fundamentalism—in both its metaphysical component and (therefore) its epistemological component—is false. Given this, and given that (II) is the only possibility given

Incomparability, we can conclude that if *Incomparability* is true then cross-domain reasons fundamentalism is false. (This will not be particularly surprising to anyone.)

Now consider the case where *Comparability Without Determination* is true. Here the reasoning will be very similar to that in the case where *Incomparability* is true, so I'll be brief. Here, as with *Incomparability*, we have two options: either (I) there is no cross-domain overriding—an option we have to reject given our assumption of the overriding view—or (II) cross-domain overriding is to be explained in terms of something other than the relative weights of conflicting reasons. And as we've seen, if (II) is true, then cross-domain reasons fundamentalism in both its components is false. And so if *Comparability Without Determination* is true, then cross-domain reasons fundamentalism is false.

7.2.1.2. *Comparability With Determination*

If *Comparability With Determination* is true, then the metaphysical component of cross-domain reasons fundamentalism is definitely true. But I want to argue that, even in that case, the epistemological component is still false. I begin with the premise that

Overriding is an antisymmetric relation

If O_A overrides O_B , then O_B doesn't override O_A .

This premise is highly plausible. When we say that Monica's moral obligation overrides her prudential obligation (for example) we mean *at least* this: Monica's moral obligation enters into some kind of conflict with her prudential obligation, and her moral obligation emerges victorious from this conflict, and her prudential obligation emerges defeated. Such a victory for morality would be impossible if it is shared equally on both sides of the conflict. In this sort of competition, if both sides win the prize, then neither side can rightfully be said to have won or lost.

My next step is to make two observations: From within the prudential standpoint, it looks as if Monica's prudential reason outweighs Monica's moral reason; and from within the moral standpoint, it looks as if Monica's moral reason outweighs Monica's prudential reason. These claims, I believe, are well-supported by experiences familiar to anyone who has occupied these normative standpoints.

Let me draw that phenomenological claim out a bit. My claim is that, for instance, anyone who knows what it is like to consider matters from the prudential standpoint knows that from within that standpoint it appears that—for example—Monica has most reason to do precisely what she has most prudential reason to do, which is to say that it appears that she

has most reason to invest rather than to donate. Another way to say this is to say that from the prudential standpoint, it seems utterly irrational for Monica to donate the money. If she donates the money, then that money will just be gone—it'll do no good for her whatsoever. The fact that donating the money would help others (even to the point of saving many lives) either does not register as a reason at all from the prudential standpoint or, if it registers as a reason, it registers as a reason that is clearly outweighed by Monica's reason to invest. A financial advisor—whose job is to view Monica's situation from the prudential standpoint and to offer advice from within that standpoint—would surely concur. The financial advisor's verdict would be that Monica has all the reason in the world to invest and little to no reason to donate.⁵

Further, my phenomenological claim is that something similar is true of the moral standpoint: From within the moral standpoint, it appears that Monica has most reason to do precisely what she has most moral reason to do—that is, she has most reason to donate rather than to invest. This, we can suppose, is what a competent moral advisor would say to Monica. And, further still, I want to make a general claim: that from within any given standpoint X, it appears that any given agent has most reason to do whatever she has most X-reason to do.

Now, what this means is that we cannot accept the following principle:

Outweighing principle

As long as it appears to us that the reasons underlying O_A outweigh the reasons underlying O_B , then we are justified in concluding that O_A overrides O_B .

The reason we cannot accept the outweighing principle is that—given my phenomenological claims above—the outweighing principle would typically justify us in believing that any two conflicting obligations from different domains override one another. In the case of Monica, for example, by switching back and forth between the moral standpoint and the prudential

⁵ One might object that from the prudential standpoint, non-prudential (e.g., moral) reasons do appear to have significant weight, or would appear to have significant weight if they were given consideration—but while one is occupying the prudential standpoint one is merely *bracketing* or *ignoring* non-prudential reasons. In response, consider the fact that if Monica were to donate then she will save many others' lives. This fact is a weighty moral reason. Yet, while Monica is occupying the prudential standpoint, she can focus her attention squarely and directly on this fact; she needn't ignore it. Indeed, she likely needs to take account of this fact in order to thoroughly consider her situation from the prudential standpoint, because she needs to determine whether she could personally benefit in any way by saving others' lives. So it is not the case that occupying the prudential standpoint means ignoring moral reasons. Thanks to an anonymous referee for pushing me on this point.

standpoint and applying the Outweighing Principle from within each standpoint, we'd quickly gain license to conclude that Monica's moral obligation both outweighs and is outweighed by her prudential obligation. But we should not be able to conclude this, given the premise that overriding is an antisymmetric relation. So we should reject the outweighing principle.

The upshot of this is that, even if the metaphysical component of cross-domain reasons fundamentalism is true, the epistemological component is false (because the epistemological component requires the outweighing principle—if knowledge requires justification). Summing up: Even when we assume *Comparability With Determination*, we should still conclude that the Outweighing Principle is false; and if the Outweighing Principle is false, then the epistemological component of cross-domain reasons fundamentalism is false; so, even if *Comparability With Determination* is true, the epistemological component of cross-domain reasons fundamentalism is false.

At this point I've made the following argument (strewn between Sections 7.2.1.1 and 7.2.1.2):

The argument against the epistemological component of reasons fundamentalism

- (A) Either *Incomparability*, *Comparability Without Determination*, or *Comparability With Determination* is true.
- (B) If either *Incomparability*, *Comparability Without Determination*, or *Comparability With Determination* is true, then the epistemological component of reasons fundamentalism is false.
- (C) Therefore, the epistemological component of reasons fundamentalism is false.

I will now go on to draw out some of the interesting implications of the conclusion of this argument.

7.2.1.3. Novel Higher-Order Domains

I have argued that the epistemological component of reasons fundamentalism is false. The upshot of this is that, in order to determine that (for example) Monica's moral obligation overrides her prudential obligation, it is not enough to observe from within the moral standpoint that Monica's moral reason outweighs her prudential reason. Some further step is required. But what is that further step?

If the needed insight is unavailable from within the moral standpoint, then it would stand to reason that we must shift to some other standpoint. But this other standpoint cannot be the prudential standpoint. That's a non-starter. From within the prudential standpoint it appears that Monica's

prudential reason outweighs her moral reason, and—no matter what the relation between reasons and cross-domain overriding turns out to be—this cannot be a basis for the conclusion that Monica's moral obligation overrides her prudential obligation.

So, if we can determine that Monica's moral obligation overrides her prudential obligation, it seems that this must be determinable only by occupying some third standpoint—call it Standpoint D_1 . From within Standpoint D_1 , we may suppose, one gets an external view of the conflict between the moral and the prudential—and, somehow, it is from that external view that one can justifiably conclude that (in Monica's case at least) the moral obligation overrides the prudential one but not vice versa.

A natural question at this point is: What is Standpoint D_1 ? Many possibilities arise here: Maybe Standpoint D_1 is *the rational standpoint*, or is in some way an *impartial* standpoint, or the standpoint that results from reflective equilibrium reasoning, or (borrowing a possibility from Sidgwick 1874) *the point of view of the universe*, or (borrowing a possibility from Nagel 1970) *the view from nowhere*, or (borrowing a possibility from Foot 1978) is *unscripted*, or (borrowing a possibility from McLeod 2001) is the domain of *just plain ought*. Of course, we have to be careful here because some of these kinds of characterizations have variously been offered of the moral standpoint itself, and I have already argued that Standpoint D_1 isn't the moral standpoint. In any event, for my purpose, I don't require a worked-out view about what Standpoint D_1 is or how we should refer to it. What I require is only the negative characterization of Standpoint D_1 : that Standpoint D_1 isn't the moral standpoint, or the prudential standpoint, or any of the other standpoints that may be in conflict at the zeroth-order level. This means that Standpoint D_1 is, in the terms I introduced above, a novel higher-order standpoint.⁶

Now I have argued that if it is determinable that (for example) Monica's moral obligation overrides her prudential obligation, then there must be a novel higher-order standpoint—Standpoint D_1 —from which this determination can be made. This isn't *quite* the same as the conclusion that there in fact is at least one novel higher-order standpoint. After all, it's possible that we simply cannot determine that Monica's moral obligation overrides her prudential obligation. Perhaps such things are unknowable.

⁶ Standpoint D_1 could be a standpoint that takes account of both prudential and moral reasons—call it *prudence-plus-morality*. Even then, Standpoint D_1 would not be the moral standpoint and wouldn't be the prudential standpoint, for the simple reason that *prudence-plus-morality* takes account of more reasons than either the prudential standpoint or the moral standpoint. So *prudence-plus-morality* would still count as *novel* in my terms.

But we've assumed the overriding view—so we've assumed that cross-domain overriding occurs. And if cross-domain overriding does indeed occur, then—if we aren't extreme skeptics—we should think that facts about cross-domain overriding are at least in principle knowable. Richard Rowland (2016) persuasively argues for a general constraint according to which any given normative fact must be knowable by some possible agent; and this general constraint is all that I require for the argument here (given that facts about overriding are normative facts). Further, my argument here is precisely that facts about cross-domain overriding are knowable only by adopting some novel higher-order standpoint. So we should now conclude that there is at least one novel higher-order standpoint, which I've called Standpoint D_1 .

Now recall the relation between standpoints and domains. I claimed above, in introducing the notion of a standpoint, that a given standpoint counts as the X-standpoint because obligations of the X-domain are apprehensible from that standpoint. Given this claim we should say that Standpoint D_1 is associated with a normative domain—and because Standpoint D_1 is a novel standpoint, the associated domain is novel also. This associated domain—which we may call D_1 —must (like any normative domain) contain deontic objects and reasons. Presumably, one of the deontic objects that belongs to this domain is an obligation that points toward the overriding domain at the zeroth-order level—in the case of Monica, this would be the moral domain.

So the picture that is now emerging is, as applied to the case of Monica, something like this: When we take up the prudential point of view, we see that Monica must invest; when we take up the moral point of view, we see that Monica must donate; to adjudicate this conflict, we take up Standpoint D_1 (which we may say is the rational standpoint, or the point of view of the universe, or the unsubscripted standpoint, or whatever); and once we do this—that is, once we take up Standpoint D_1 —then we see that Monica must take up the moral standpoint rather than the prudential standpoint in order to decide what to do (i.e., she must act in line with her moral obligation rather than her prudential one).

7.2.2. From Novel Higher-Order Domains to Infinitism about Cross-Domain Conflict

I have thus far argued that *in order to know* that (e.g.) a moral obligation overrides a prudential obligation, we have to be able to occupy a higher-order standpoint, Standpoint D_1 , and to see, from that standpoint, that the agent has an obligation to act in line with her moral obligation rather than

her prudential obligation. Given the connection between standpoints and domains, this means that—in general—in order to know that O_A overrides O_B we have to know that there is an obligation belonging to a higher-order domain that requires the agent to act in line with O_A rather than O_B .

This is an epistemic claim, not a metaphysical one. That is, it is not a claim about *what it is* for a moral obligation to override a prudential obligation. But this epistemic claim does support a further metaphysical claim, or so I will now argue (Section 7.2.2.1). After that, I will show how this metaphysical claim supports the argument for infinitism about cross-domain conflict (Section 7.2.2.2).

7.2.2.1. *Higher-Order Domains Ground Cross-Domain Overriding*

I want to argue that the epistemic claim that I've been defending supports the following view:

Higher-order domains ground cross-domain overriding

Suppose obligation O_A belongs to normative domain D_A and obligation O_B belong to normative domain D_B (where D_A and D_B are different domains). If O_A overrides O_B , this is *at least in part because* there is a higher-order domain, D_1 , which contains an obligation, O_1 , that requires the agent to act in line with O_A rather than O_B .

My argument for this claim is a kind of inference to the best explanation. I've argued that in order to know that O_A overrides O_B we have to know that there is an obligation belonging to a higher-order domain that requires the agent to act in line with O_A rather than O_B . The most straightforwardly appealing explanation for this epistemic constraint, I suggest, would be that the obligation belonging to the higher-order domain is at least part of what *makes it the case* that O_A overrides O_B .

Here's an analogy to support this inference. Suppose that I'm a private in an army. Sergeant_A has just given me an order to (say) march twenty miles west. But a different sergeant, Sergeant_B, has given me an incompatible order to (say) march twenty miles east. Suppose I am of the view that each sergeant's order creates a certain kind of obligation for me. So I am obligated to march twenty miles west and simultaneously obligated to march twenty miles east. I want to know which of these obligations overrides the other. Suppose I come to the view that I cannot know which obligation overrides the other without consulting someone who is further up the hierarchy than either of the sergeants—say, a lieutenant. If the lieutenant says that I ought to go with Sergeant_A's order, then I'm in a position to know that Sergeant_A's order overrides Sergeant_B's order; if the lieutenant says that I ought to go with Sergeant_B's order, then I'm in a position to know that Sergeant_B's order overrides Sergeant_A's order. Well, if I come to that epistemic view, then it

would be natural to think that the lieutenant's say-so is also at least part of what makes it the case that one sergeant's order overrides the other. In this example, then, an epistemic claim about how I can know that one order overrides another provides support for a metaphysical claim about what makes it the case that one order overrides the other.

I do not want to claim that in *every* case, if I cannot know some fact F_1 without knowing some other fact F_2 , I am entitled to conclude that F_2 is part of what makes F_1 the case. That principle is clearly false. For example, perhaps I cannot know that the temperature is fifty degrees without knowing that the thermometer reads fifty degrees, but this does not show that the reading on the thermometer is any part of what makes it the case that the temperature is fifty degrees.

However, it is the case that in general, if I cannot know some fact F_1 without knowing some other fact F_2 , this epistemic constraint calls for some sort of an explanation. And in most cases there will only be three sorts of explanations to consider: either F_2 brings F_1 about; or F_1 brings F_2 about; or some third factor brings F_1 and F_2 about. Each of these possibilities could, in principle, explain why I cannot know some fact F_1 without knowing some other fact F_2 .

In the present case, these three options would be:

- (i) The fact that [O_A overrides O_B] brings it about that [an obligation belonging to a higher-order domain requires the agent to act in line with O_A rather than O_B].
- (ii) The fact that [an obligation belonging to a higher-order domain requires the agent to act in line with O_A rather than O_B] brings it about that [O_A overrides O_B]. (This is the claim that I'm defending.)
- (iii) Some third factor brings about both that [O_A overrides O_B] and that [an obligation belonging to a higher-order domain requires the agent to act in line with O_A rather than O_B].

Each of these three possibilities could in principle explain why we need to know that [an obligation belonging to a higher-order domain requires the agent to act in line with O_A rather than O_B] in order to know that [O_A overrides O_B]. So, in order to defend (ii) (my preferred explanation) I need to say something against (i) and (iii).

(i) and (iii) have a crucial similarity. If we accept either (i) or (iii), then we are left without an explanation of the fact that O_A overrides O_B . So, if we accept either (i) or (iii), then we are suddenly in need of something to say about *why* O_A overrides O_B .

Perhaps the relative weights of reasons could explain this. On that view, the reasons underlying O_A are weightier than the reasons underlying O_B , and *this*—rather than any facts about higher-order domains—is why O_A overrides O_B . But this explanation isn't satisfactory.

Here's why. I argued in Section 7.2.1.2 that from within the moral standpoint, moral reasons seem to outweigh prudential reasons, whereas from within the prudential standpoint, prudential reasons seem to outweigh moral reasons. This suggests that reasons can be measured on different scales (as it were) and that different scales give different results. On the moral scale, moral reasons register as heaviest; on the prudential scale, prudential reasons register as heaviest. And it is unclear to me how defenders of either (i) or (iii) can identify a basis on which one or the other of these two scales is rightly regarded as uniquely correct or as more correct than the other. We may suppose that each scale measures some quantifiable attribute of reasons; each measures something that could be called 'weight.' But it is unclear how defenders of either (i) or (iii) could produce a principled reason why one of these quantifiable attributes rather than the other deserves to be regarded as the *true* weight of reasons.

Of course, at this point, in order to explain why the quantifiable attribute measured by the moral scale (or the prudential scale, or whatever scale) is the one that is correctly regarded as the true weight of reasons, proponents of (i) or (iii) might want to point to a third, higher-order normative standpoint as an arbiter between the contradictory deliverances of the diverging scales. But if we point to a third higher-order standpoint as arbiter, then the view begins to bear an uncanny resemblance to (ii).

So proponents of (i) or (iii) will have difficulty explaining why O_A overrides O_B in terms of the relative weights of reasons, at least not without allowing their view to shade into some version of (ii). Can proponents of (i) or (iii) come up with some alternative, superior explanation, one that doesn't appeal to the relative weights of reasons? I don't see what that alternative explanation might be. By contrast, (ii) is precisely an explanation of why O_A overrides O_B . This is why I think (ii) is to be preferred over (i) or (iii). And (ii) implies the claim that higher-order domains ground cross-domain overriding.

7.2.2.2. *Authority and the Beginning of the Regress*

Recall the argument for infinitism about cross-domain conflict. The first premise is just a definition:

- (1) O_A overrides O_B iff O_A belongs to D_A and O_B belongs to D_B and D_A trumps D_B .

The second premise is:

- (2) If D_A trumps D_B , this is because there is a novel higher-order domain, D_1 , that explains why D_A trumps D_B .

The above premise follows directly from the argument I gave in Section 7.2.1. The next premise in the argument is as follows:

- (3) D_1 explains why D_A trumps D_B only if D_1 has the requisite authority.

Intuitively, not just any domain can explain why D_A trumps D_B . To illustrate this, imagine a first-order analog of the esthetic domain. Suppose there is a standpoint from which we can assess the beauty (or other esthetic properties) of different zeroth-order domains. From this standpoint, suppose, it appears that the correct zeroth-order domain—that is, the zeroth-order domain that should be used to guide one's actions—is whatever happens to be the most beautiful domain. And suppose that D_A just happens to be considerably more beautiful than D_B . Then, from this first-order esthetic standpoint, it would appear that one ought to act in line with the obligations of D_A rather than those of D_B ; in other words, the domain associated with that standpoint would contain an obligation to act in line with D_A rather than D_B . But this would not mean that D_A *genuinely* trumps D_B . For it seems clear that we should not select our zeroth-order standpoint on esthetic grounds. The first-order esthetic domain that I am discussing here thus lacks what I'll call the *authority* required in order to explain why D_A trumps D_B .

The next step in the argument is:

- (4) D_1 has the requisite authority only if D_1 trumps *at least some of* its same-order competitors D_1^* , D_1^{**} , D_1^{***} ,

To defend this premise, I have to first explain what I mean by a *same-order competitor*. Imagine that from Standpoint D_1 , it appears that one should act in line with D_A , but from Standpoint D_1^* , it appears that one should act in line with D_B . Then D_1^* is a same-order competitor of D_1 . D_1^* and D_1 are same-order competitors because they contain conflicting verdicts about domains one level below them. The thought behind (4) is that if D_1 and D_1^* contain conflicting verdicts in this way, then D_1 cannot have the sort of authority mentioned in (3) unless D_1 trumps D_1^* .

To see the plausibility of this, let's once again consider the example of the army private. In that example, Sergeant_A has just given me (the private) an order to march twenty miles west, and a different sergeant, Sergeant_B, has given me an incompatible order to march twenty miles east. To resolve this conflict, I go to the rank above sergeant—the rank of lieutenant. I find a lieutenant, Lieutenant₁, who tells me to go with Sergeant_A's order. This could conceivably explain why Sergeant_A's order overrides Sergeant_B's order (i.e., it could explain why Sergeant_A trumps Sergeant_B). But now suppose that, unfortunately, a different lieutenant, Lieutenant₁^{*}, tells me to go with Sergeant_B's order. Then in this eventuality, it seems that Lieutenant₁'s verdict cannot explain why Sergeant_A trumps Sergeant_B—*unless* it is also the case that Lieutenant₁ trumps Lieutenant₁^{*}. The thought behind (4) is that something similar needs to be true in the analogous case with regard to normative domains.

Here one might issue the following objection. There is no guarantee that Lieutenant₁ will *have* any same-rank competitors. Perhaps Lieutenant₁ is the only lieutenant. Or perhaps there are multiple lieutenants, but all of the lieutenants agree (and therefore aren't competitors): perhaps all of them tell me to heed Sergeant_A rather than Sergeant_B. In either case, we might reasonably think that Lieutenant₁ doesn't need to trump any same-rank competitors in order for her verdict to explain why Sergeant_A trumps Sergeant_B. Likewise, one might suggest, it is not a foregone conclusion that D₁ has any same-order competitors: perhaps D₁ is the only first-order domain; or perhaps all of the first-order domains agree with D₁. In that case we could imagine that D₁ has the requisite authority without trumping any competitors at all, and then (4) would be false.

However, I think it is inevitable that D₁ *will* have same-order competitors. It seems that, at every level, there should be a domain for every relevant principle. (This is not the same as to say that every domain is governed by a principle. There may be unprincipled domains. Indeed, if moral particularists are correct, morality is one such domain.) At the zeroth order, the relevant principles are action-guiding principles. Given that there are infinitely many different logically possible action-guiding principles, there should be infinitely many domains at the zeroth order. At the first order, the relevant principles are principles that guide the selection of zeroth order domains, and there are infinitely many such principles. So we should think that D₁ has infinitely many same-order competitors; and the subset of these with which D₁ conflicts will likewise be infinite. So I believe there is no way for D₁ to avoid having same-order competitors.

The fifth premise of the argument is as follows:

- (5) If D₁ trumps any of its same-order competitors, this is because there is a novel higher-order domain, D₂, that explains why D₁ trumps those same-order competitors.

This premise can be defended with a line of argument that is precisely parallel with the above defense of (2). This is where the regress begins. To explain why D₁ trumps its same-order competitors, we need D₂; in order for D₂ to do this explanatory work, D₂ needs authority; in order for D₂ to have authority, D₂ needs to trump at least some of its same-order competitors D₂^{*}, D₂^{**}, D₂^{***}, ...; in order for D₂ to trump its same-order competitors, there needs to be a novel higher-order domain D₃ to explain this; and so on.

The result: In order for O_A to override O_B, there needs to be an infinitely long chain of domains D₁, D₂, D₃, ... such that D₁ contains an obligation to act in line with the obligations of D_A, and D₂ contains an obligation to act in line with the obligations of D₁, and D₃ contains an obligation to act in

line with the obligations of D_2 , and so on, indefinitely. And that's infinitism about cross-domain conflict.

The picture of normativity that emerges from this infinitism is something like the following. We are at least sometimes confronted by a variety of normative domains that pull us in different directions. And in at least some of these cases, one domain trumps the other—for example, perhaps in a given case morality wins out over prudence. When this happens, it happens in virtue of a higher-order domain that possesses the requisite authority. We can occupy the standpoint associated with that higher-order domain by *taking a step back*, as it were, and viewing both morality and prudence from the outside, that is, from the perspective of that authoritative higher-order domain. Once we do this, we can ask the question of why this higher-order domain possesses the authority that it does—that is, we can ask why this higher-order domain (rather than any of its possible alternatives) is the domain we should rely upon in order to adjudicate the conflict between morality and prudence. And in order to address that further question, we must take a step back yet again, and view from the outside that higher-order domain alongside its alternatives. The picture given by infinitism is that we will *always* be able to take these sorts of steps backward, no matter how many steps back we take. We will never reach a point where the authority of a given domain to adjudicate lower-order domains is brutally given.

7.3. OBJECTIONS AND REPLIES

In this section I will consider a few of the many objections that may be raised against the argument that I have just laid out. Several of these objections are directed against infinitism about cross-domain conflict, rather than against the conditional claim that *if* the overriding view is true, then infinitism is true. But such objections may be taken as indirectly targeting the conditional claim, given that such objections would show that the conditional claim ought to be rejected by anyone who finds the overriding view to be plausible.

First objection: Infinitism requires that there is an infinite series of domains and an infinite series of obligations contained in those domains. But it is implausible to suppose that there are infinitely many of either of those sorts of things.

To start with, it is not implausible to suppose that there are infinitely many obligations. Indeed, it would be surprising if there weren't. For there seem to be infinitely many (real and hypothetical) cases in which agents have

to make up their minds about what to do and are obligated to do one thing rather than another.

It is true that infinitism implies that *any given agent* has infinitely many obligations. In the example we've been working with, an agent has an obligation given by D_A (this is an obligation to act in a certain way), and has an obligation given by D_1 (this is an obligation to act in line with D_A rather than D_B), and has an obligation given by D_2 (this is an obligation to comply with D_1 rather than D_1^* , D_1^{**} , D_1^{***} , ...), and so on. Perhaps it seems strange to think that *one agent* can be saddled with infinitely many obligations in this way. But the idea that one agent can have infinitely many obligations is also not particularly strange.

Suppose a given agent is obligated to give at least 4 percent of her income to charity. Then it would seem to follow trivially that she is obligated to give at least 3.999 percent of her income to charity, and is obligated to give at least 3.998 percent of her income to charity, and so on. In this way we could start with one obligation and use it to generate an infinite number of further obligations for one agent. So the very idea of one agent having infinitely many obligations isn't objectionable.

The other part of the present objection points out that infinitism requires an infinite series of domains as well as obligations. But I am not sure why that would be objectionable if infinitely many obligations isn't objectionable.

Second objection: Infinitism makes cross-domain overriding unknowable. In order to know that D_A trumps D_B , we have to know that D_1 has authority; and in order to know that D_1 has authority, we have to know that D_1 trumps its same-order competitors; and in order to know that D_1 trumps its same-order competitors, we have to know that D_2 has authority; and so on. Thus infinitism requires us to know an infinite series of these sorts of normative facts. Further, we cannot know these sorts of facts without occupying an infinite series of standpoints associated with these domains. For example, we cannot know that D_A trumps D_B without occupying the standpoint associated with D_1 ; and we cannot know that D_1 trumps its same-order competitors without occupying the standpoint associated with D_2 ; and so on. So, in order to know that D_A trumps D_B , we need to be capable of occupying an infinite series of standpoints, one after the other. It is not possible to do that (given that occupying a standpoint requires some positive amount of time; see Reisner 2015 for relevant discussion) so it is not possible to know that D_A trumps D_B .

I think it is probably true that we cannot know that D_A trumps D_B without occupying the standpoint associated with D_1 . And I think it is probably also true that we cannot know that D_1 trumps its same-order competitors without occupying the standpoint associated with D_2 . But I don't think this

has to be the case at every level. At some point, I suspect, an ordinary agent may be entitled to make an inductive inference. If the agent finds that, no matter how far back along the chain of domains she goes, she always finds that there is a domain one level up that continues the pattern she has thus far found to hold, she may be justified in inferring that the pattern will continue indefinitely. Of course, she may never be entitled to *certainty* that the pattern will continue indefinitely. But in general it is highly plausible that inductive inferences do not need to justify certainty in order to confer knowledge.

Third objection: The argument requires a flawed and even absurd account of authority. On the required account, the authority of a given domain at the zeroth order is transmitted from a higher-order domain at the first order. But the domain at the first order must first receive the authority thus transmitted from a higher-order domain at the second order, and so on, indefinitely. In this account there is no highest order domain to originate the authority that gets transmitted along the infinite series of domains; and so the authority that is being transmitted from domain to domain comes from nowhere. And if it comes from nowhere, it cannot exist at all.

It is a bit odd to suggest, as this objection does, that my argument requires authority to be *transmitted* along an infinite series of domains. This makes it sound as if authority is like a single baton being passed from runner to runner. But in my account, there isn't a single object that is passed from domain to domain. Rather, each domain has its own authority: the authority to adjudicate conflicts at the domain below it. The authority of D_1 to adjudicate conflicts at the zeroth order is a different thing (or, more exactly, it is a different token property) from the authority of D_2 to adjudicate conflicts at the first order.

The objection would be slightly improved if it were to claim that, on my account, the authority of any given domain derives from the *separate* authority of a higher-order domain. But once this is clarified, it is not clear to me that there is a serious objection here. Provided that we can have an infinite series of domains (a point I've already addressed) I don't see what is absurd or even objectionable about a picture in which each domain in such a series receives its authority from the next one in the series. After all, " D_Y bestows authority on D_X " just names a relation between D_Y and D_X , and it is possible for each element in an infinite series to bear relations to the other elements; for example, each natural number in the infinite series of natural numbers is greater than the number below it.

Fourth objection: According to this argument, the power of morality to override prudence, in a given instance, requires the existence of an infinite chain of domains in which D_1 favors morality over prudence, D_2 favors D_1

over D_1^* , D_3 favors D_2 over D_2^* , and so on. Call this *the morality-favoring chain*. But there must be multiple such infinite chains. For example, it must be that there is a chain such that D_1^* favors prudence over morality, and D_2^* favors D_1^* over D_1 , and so on. Call this *the prudence-favoring chain*. (This must be the case because, as I argued earlier, at every level, there should be a domain for every relevant principle.) If there are two of these sorts of infinite chains each pointing toward conflicting verdicts about which zeroth order domain to take up, then this seems problematic. For, if there is both a morality-favoring chain and a prudence-favoring chain, then it would follow that morality trumps and is trumped by prudence; this would be objectionable because (as we've seen) overriding is an antisymmetric relation.

However, at no point in the argument have I laid out sufficient conditions for overriding, for authority, or for any of the other normative properties that I've been discussing here. So my argument does not commit me to say that the existence of a morality-favoring chain in a given case is sufficient to show that morality trumps prudence in that case and likewise does not commit me to say that the existence of a prudence-favoring chain in a given case is sufficient to show that prudence trumps morality in that case.

My argument is consistent with the possibility that, in order for morality to trump prudence, there must be a morality-favoring chain *and* there must be some further factor present. Here is one attractive possibility along those lines. Suppose that there is a property of *normative oomph*. (For articulation of the hard-to-articulate notion of normative oomph, see Joyce 2006; Copp 2015: 142.) It may be that a given domain D_X has authority in a given case iff (i) D_X trumps at least some of its same-order competitors, *and* (ii) D_X is, as we might say, *oomphy* (i.e., the normative objects contained in D_X have the property of oomph). In that case, we could show that, for example, morality trumps prudence iff there is a morality-favoring chain of domains *and* every domain in that chain is oomphy. And this, in turn, would mean that even if there is both a morality-favoring chain and a prudence-favoring chain in a given case, it might still be the case that morality trumps prudence but not vice versa, or that prudence trumps morality but not vice versa. In this way, my argument can respect the constraint that overriding is an antisymmetric relation.

Here it may be asked: What if, in a given case, there is both a morality-favoring chain and a prudence-favoring chain, and every domain in *each* of those chains is oomphy? Then, in order to respect the antisymmetry constraint, we'd need to modify the account again. For example, we may say that a given domain D_X has authority in a given case iff (i) D_X trumps at least some of its same-order competitors, and (ii) D_X is oomphy, and (iii) none of D_X 's same-order competitors are oomphy. This is just one

possibility along the present lines; there are many others. But this possibility should suffice to show that I can allow the possibility of multiple infinite chains of domains in a given case without falling afoul of the antisymmetry constraint.

At this point one might suggest that if the notion of oomph is brought into the picture, we can rely on this notion to explain why one domain trumps another, and we do not need to consider higher-order domains at all. But this suggestion faces a dilemma.

Either oomph is defined in terms of overridingness, such that a given obligation has oomph iff the obligation overrides all obligations with which it conflicts, or it isn't. A cost of defining oomph in terms of overridingness is that it will imply that any two same-order competitor domains cannot both be oomphy. Further, and more importantly in the present context, if oomph is defined in that way then we will not be able to rely on oomph to explain overridingness (because oomph will simply *be* overridingness).

Alternatively, if oomph is not defined as overridingness, then it is defined in some other way or not defined at all. In either case oomph is unlikely to be sufficient on its own to explain why one domain trumps another. Suppose that the moral domain is oomphy and that the legal domain is not oomphy. Then given that oomph is not defined as overridingness, it is still an open question whether morality trumps legality. Indeed, the possibility that a non-oomphy domain could trump an oomphy one is not a mere logical possibility. For instance, if an oomphy higher-order domain favors legality over morality, then it may be plausible that legality trumps morality even on the supposition that morality is oomphy and legality isn't. If that's right, then it shows that oomph alone is not sufficient to explain why one domain trumps another.⁷

References

- Copp, David (1997). "The Ring of Gyges: Overridingness and the Unity of Reason," *Social Philosophy and Policy* 14 (1): 86–106.
- Copp, David (2015). "Rationality and Moral Authority," in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics, Volume 10*. Oxford: Oxford University Press, 134–59.

⁷ For extremely valuable feedback, advice, and encouragement, I'm exceedingly grateful to Thomas Carson, Guy Fletcher, Rebecca Harrison, Richard Rowland, Russ Shafer-Landau, Rob Streiffer, and many individuals who were in the audience at the Metaethics Workshop in Madison.

- Copp, David (ms.). "Normative Pluralism and Skepticism about Ought Simpliciter."
- Crisp, Roger (ms.). "Partiality and the Dualism of Practical Reason."
- Cuneo, Terence and Shafer-Landau, Russ (2014). "Moral Fixed Points," *Philosophical Studies* 171 (3): 399–443.
- Dorsey, Dale (2016). *The Limits of Moral Authority*. Oxford: Oxford University Press.
- Enoch, David (2019). "How Principles Ground," in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics, Volume 14*. Oxford: Oxford University Press, 1–22.
- Foot, Philippa (1978). *Virtues and Vices: And Other Essays in Moral Philosophy*. Oxford: Blackwell and Berkeley: University of California Press.
- Hare, R. M. (1981). *Moral Thinking: Its Levels, Method, and Point*. Oxford: Oxford University Press.
- Joyce, Richard (2006). *The Evolution of Morality*. Cambridge, MA: MIT Press.
- Kearns, Stephen and Star, Daniel (2008). "Reasons: Explanations or Evidence?" *Ethics* 119 (1): 31–56.
- Klein, Peter (1998). "Foundationalism and the Infinite Regress of Reasons," *Philosophy and Phenomenological Research* 58 (4): 919–25.
- McElwee, Brian (2017). "Supererogation Across Normative Domains," *Australasian Journal of Philosophy* 95 (3): 505–16.
- McLeod, Owen (2001). "Just Plain 'Ought,'" *The Journal of Ethics* 5 (4): 269–91.
- McPherson, Tristram (2018). "Authoritatively Normative Concepts," in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics, Volume 13*. Oxford: Oxford University Press, 253–77.
- Nagel, Thomas (1970). *The Possibility of Altruism*. Princeton, NJ: Princeton University Press.
- Parfit, Derek (2011). *On What Matters: Volume 1*. Oxford: Oxford University Press.
- Reisner, Andrew (2015). "Fittingness, Value, and Trans-World Attitudes," *Pacific Philosophical Quarterly* 65 (260): 464–85.
- Rowland, Richard (2016). "The Significance of Significant Fundamental Moral Disagreement," *Nous* 51 (4): 802–31.
- Sidgwick, Henry (1874). *The Methods of Ethics*. London: Macmillan.
- Tiffany, Evan (2007). "Deflationary Normative Pluralism," *Canadian Journal of Philosophy* 37 (33): 231–62.
- Wolf, Susan (2015). *The Variety of Values*. Oxford: Oxford University Press.